

APPENDIX A

Hitachi Freedom Storage

Hitachi Freedom Storage™ Software

Solutions Guide

Best-of-Breed Software for
Centralized Storage Management


Storage Area Management Solutions

Business Continuity Solutions

Backup and Recovery Solutions

Data Movement Solutions

Professional Services and Support



Best-of-Breed Software for Centralized Storage Management

HITACHI
DATA SYSTEMS

Table of Contents

Table of Figures	viii
Table of Tables	ix
Chapter 1 – <i>Best-of-Breed</i> Software for Centralized Storage Management	1
Hitachi TrueNorth – A Vision for Open Information Management	1
The Business Value of Hitachi Freedom Storage Software Solutions	2
Storage Area Management Solutions	2
Performance Enhancement Solutions	3
Business Continuity Solutions	3
Backup and Recovery Solutions	4
Data Movement Solutions	5
Other Information Sources Available from Hitachi Data Systems	5
Chapter 2 – Storage Area Management and Performance Enhancement Solutions	7
HiCommand: Hitachi's Framework for Open Storage Management	7
No Single Vendor Can Meet All Storage Requirements for Today's Enterprise.	7
Open, Standards-based Approach	8
HiCommand Allows Global Systems Management of Multi-vendor Storage Systems	9
The Components of HiCommand Device Manager Software	9
How Hitachi Resource Manager Supports Business Objectives	10
Hitachi Resource Manager is Modularly Designed	11
1. Remote Console - Storage Navigator	12
2. Graph-Track Software Takes the Guesswork Out of Performance Management	12
3. SNMP Support Ensures Compliance with All Major Systems Management Platforms	13
4. Sun Java Agent	14
5. Host Storage Domains	14
6. Hitachi LUN Manager and LUN Expansion Manager Software Simplify Configuration Management and Reduce Staffing Costs	15
7. SANtinel Software	15
8. Virtual LUN and LUN Expansion Software	16
9. Virtual Logical Volume Image VLVI Manager for S/390 Environments Allows Very Large to Very Small Virtual Volume Configuration	17
10. Hitachi FlashAccess Software Allows Data to Be "Locked and Unlocked" in Cache on-the-Fly ..	17
11. Cache Manager Host Agent	18

Hitachi Performance Enhancement Solutions	18
Hitachi Performance Maximizer Package	18
1. Performance Monitor Software	18
2. CruiseControl Software	18
3. Priority Access Software	19
Transaction Processing Facility/Multi-path Lock Facility (TPF/MPLF)	19
Hitachi Parallel Access Volume (HPAV) and Multiple Allegiance (MA)	20
VERITAS Software Simplifies Management and Increases Availability	20
1. VERITAS Volume Manager	20
2. VERITAS File System	21
3. VERITAS Storage Administrator Improves System Manageability	21
4. VERITAS SANPoint Control	22
Chapter 3 – Business Continuity Solutions	23
Increased IT Service Levels in Availability and Performance through Business Continuity Software Solutions and Services	23
The “Scale of 9s” Is an Accepted Measure of Computer System Availability	24
The Importance of the Continuous Business Paradigm to Nonstop Operations	24
Hitachi Business Continuity Software Solutions and Services Provide a World-class Selection of Risk Reduction Alternatives for the Enterprise	27
1. Remote Copy	27
2. PiT Copy	28
3. Data Migration	28
Hitachi TrueCopy Solutions	28
TrueCopy Basic (Synchronous)	29
TrueCopy (Asynchronous)	31
NanoCopy	35
Hitachi ShadowImage	38
ShadowImage Provides the Only Online Backup Solution for Multi-terabyte, Mission-critical Oracle Databases	39
Hitachi Extended Remote Copy (HXRC)	40
Hitachi Dynamic Link Manager Provides Path Failover and Load Balancing	40
Alternate Pathing, Host Failover, and Parallel Database Clusters	41
Alternate Pathing Middleware Switches the I/O Load in the Event of Path Failure	41
Host Failover	41

Parallel Database Clustering	42
The Hi-Track Remote Maintenance Tool Ensures the Accuracy of Business Continuity Solutions	42
Chapter 4 – Backup and Recovery Solutions	43
Overview of Hitachi Freedom Storage Backup and Recovery Software Solutions	43
Backup/Recovery Is Critical to High-availability Computing	43
Six Alternative Techniques for Backup and Restore	44
1. Volume-level Backup/Restore	44
2. Hitachi Multiplatform Backup/Restore Provides Mainframe Benefits	44
3. The Tantia Technology HARBOR File-level Backup and Restore Agent	46
4. Database and Application-level Backup	46
5. Application-level Backup/Restore	47
6. Hitachi Freedom Data Networks	47
Industry Alliances are Key to Hitachi Freedom Data Networks Solutions	48
Tantia Technologies HARBOR Backup/Recovery and HYPERTape	48
VERITAS NetBackup Offers Backup and Recovery Solutions for Heterogeneous Environments	48
Chapter 5 – Data Movement Solutions	49
Hitachi RapidXchange Software Enhances Decision Support Productivity while Improving Operations Efficiency	49
File Access Library API	50
File Conversion Utility	51
Tantia Technologies HARBOR File Transfer Increases Speed and Automates Performance of Large File Transfers	51
Hitachi e-Copy Provides LANfree and Serverfree Data Movement	52
Chapter 6 – Professional Services and Support	55
Hitachi Data Systems Is Consistently Ranked Number One in the Industry.....	55
Professional Services Overview.....	55
Data Protection Services	55
Data Protection Services Lab	56
Remote Copy Assessment and Implementation Service	56
NanoCopy Implementation Service	56
Continuous Availability Service	56
Glossary of Terms	57

Table of Figures

Figure 1	HiCommand Management Framework allows management of multi-vendor hardware and software from the customer's platform of choice.	3
Figure 2	HiCommand Management Framework allows global systems management of Hitachi Freedom Storage hardware and software solutions.	4
Figure 3	The HiCommand Management Framework is based on CIM and SOAP industry standards. . .	8
Figure 4	The HiCommand GUI browser is intuitive and easy to use.....	10
Figure 5	The GUI interface on Hitachi Graph-Track software simplifies performance management. . .	13
Figure 6	Host Storage Domains reduce costs through fewer physical connections.	14
Figure 7	SANtinel software uses pre-defined zones to protect data from unauthorized access.	16
Figure 8	Priority Access software allows administrators to allocate bandwidth in line with quality of service.	19
Figure 9	Hitachi Data Systems changes the paradigm of high-availability computing from business continuity to continuous business.	25
Figure 10	TrueCopy basic for open systems environments.	30
Figure 11	TrueCopy basic for S/390 environments.	30
Figure 12	TrueCopy asynchronous extension for open systems.	32
Figure 13	TrueCopy asynchronous extension for S/390 Environments.	33
Figure 14	Not using individually timestamped I/Os can force an unplanned recovery situation.	34
Figure 15	TrueCopy software is compatible with IBM GDPS clusters.	34
Figure 16	For disaster recovery, NanoCopy triplets ensure there is always a time-consistent set of volumes in the remote system that is not being updated.	37
Figure 17	ShadowImage software supports up to nine image copies in open systems and three image copies in S/390 from a single-source volume. The image copies can then be split for other uses such as backup or application testing.	39
Figure 18	ShadowImage software can provide high-speed backup of open database systems such as Oracle.	39
Figure 19	Dynamic Link Manager automatically provides path failover and load balancing for open systems.	41
Figure 20	Multiplatform Backup/Restore software allows UNIX or Windows NT/2000 volumes to be backed up to tape using FDR and mainframe procedures.	45
Figure 21	A major benefit of backup/recovery with the HARBOR Agent for Lightning 9900 V Series systems is the ability to restore only single files rather than restoring entire data volumes. . . .	46
Figure 22	RapidXchange software quickly converts mainframe EBCDIC files to open systems ACSII files for nondisruptive population of data marts and data warehouses.	50
Figure 23	HARBOR File Transfer runs on RapidXchange software and transfers large data files at channel speeds in either direction between open systems and mainframe hosts.	51
Figure 24	Working with the VERITAS NetBackup series of products, the Hitachi e-Copy feature provides LANfree and serverfree data movement	53

Table of Tables

Table 1	Summary of Hitachi Freedom Storage software products.	2
Table 2	Summary of Hitachi Resource Manager solutions and the products they support.	11
Table 3	Summary of Hitachi Performance Enhancement Software Solutions and the products they support.	18
Table 4	The “Scale of 9s” is an accepted measure of availability.	24
Table 5	Through the Six-Sigma approach, Hitachi Data Systems attacks both planned and unplanned downtime.	26
Table 6	Hitachi Data Systems Business Continuity Solutions.	27
Table 7	Hitachi Freedom Storage copy products provide remote copy, data duplication, and data migration capabilities.	28
Table 8	Summary of Hitachi TrueCopy solutions and the products they support.	29
Table 9	Summary of Hitachi ShadowImage solutions and the products they support.	38
Table 10	Summary of Hitachi Freedom Storage backup and recovery solutions and the products they support.	43
Table 11	Data Movement Solutions provided by Hitachi Data Systems.	49

Best-of-Breed Software for Centralized Storage Management

1

Hitachi TrueNorth™ – A Vision for Open Information Management

As the information sources driving businesses expand, organizations must manage the escalation in the physical quantity of storage, the reliance on universal access to information, the growing complexity of their storage environments, and the adoption of emerging technologies. Few businesses have unlimited resources or time to address these challenges.

Today, companies need to consider a new storage management strategy based on high-performance, intelligent systems and sophisticated software that enable total manageability of existing investments, while maximizing uptime and lowering the total cost of ownership (TCO). Hitachi TrueNorth™ provides a vision and strategic direction, along with an associated product road map, for an open information management framework that meets these challenges. With TrueNorth, we view the future storage infrastructure as a synergistic combination of management software and powerful, intelligent storage systems that will deliver always-on, set-and-forget management, complemented by flexible provisioning capabilities.

As first steps to achieving the TrueNorth vision, Hitachi Data Systems is working with its partners to expand the HiCommand™ Management Framework. Open and collaborative by design, the HiCommand Management Framework adheres to accepted industry standards and provides policy-based automation tools. To power this software we introduce Hitachi Freedom Storage™ Lightning 9900™ V Series systems, our second-generation systems based on Hi-Star™ switched architecture, delivering the industry's highest bandwidth, availability, and scalability. The HiCommand Management Framework, Lightning 9900 V Series, and the new, modular Hitachi Freedom Storage Thunder 9500™ V Series systems, combined with Hitachi Freedom Storage software solutions, are central to the TrueNorth vision and will allow organizations to *simplify, protect, and optimize* their information infrastructures, reducing TCO and ensuring a more rapid return on investment (ROI).

Customers and analysts regard the Lightning 9900 V Series, Thunder 9500 V Series, and Hitachi Freedom Storage Thunder 9200™ storage systems as the most advanced products in their markets. Recently, Hitachi has also won the respect of many analysts for having taken the leadership role in advanced function storage software that supports all Hitachi Freedom Storage products. The *Hitachi Freedom Storage Software Solutions Guide* discusses the advantages of this suite of leading-edge software products that enhances the Lightning 9900 V Series, the Thunder 9500 V Series, and Thunder 9200 storage systems. This guide is available from Hitachi Data Systems in PDF form at <http://www.hds.com>.

Hitachi Freedom Storage Software Solutions support an enterprise's strategic goal of accessing any information, on any computer, located anywhere, at any time. The many advanced functions available on Lightning 9900 V Series, Thunder 9500 V Series, and Thunder 9200 storage systems are initiated, managed, and controlled through these powerful software programs. Table 1 summarizes how Hitachi Freedom Storage Software Solutions apply to open systems and mainframe environments.

The TrueNorth vision and strategic direction is for policy-based automated management of storage investments, continuous availability, and lower cost.

Hitachi Freedom Storage Software Solutions are the most advanced available.

*Table 1 – Summary of
Hitachi Freedom Storage
software products.*

Software Suite	Product Support				OS Support	
	Thunder 9531V™ Thunder 9532V™ Thunder 9533V™	Thunder 9570V™	Lightning 9900 Series	Lightning 9900 V Series	Open Systems	IBM® S/390®
Storage Area Management Suite						
HiCommand™ Device Manager	✓	✓	✓	✓	✓	
HiCommand Tuning Manager	✓	✓	✓	✓	✓	
Hitachi Resource Manager™	✓	✓	✓	✓	✓	✓
Hitachi Performance Enhancement Suite						
Hitachi Parallel Access Volume			✓	✓		✓
Hitachi Performance Maximizer				✓	✓	✓
- Hitachi CruiseControl™			✓	✓	✓	✓
- Hitachi Priority Access				✓	✓	✓
- Hitachi Performance Monitor				✓	✓	✓
Hitachi Dynamic Link Manager™	✓	✓	✓	✓	✓	
Hitachi Resource Manager™	✓	✓	✓	✓	✓	✓
VERITAS® Volume, File and SAN Management™	✓	✓	✓	✓	✓	
Hitachi Business Continuity Suite						
Hitachi TrueCopy™		✓	✓	✓	✓	✓
Hitachi ShadowImage™	✓	✓	✓	✓	✓	✓
NanoCopy™			✓	✓		✓
Hitachi Extended Remote Copy (HXRC)			✓	✓		✓
Hi-Track®	✓	✓	✓	✓		✓
Hitachi Dynamic Link Manager	✓	✓	✓	✓	✓	
Hitachi Alternate Pathing and Host Failover	✓	✓	✓	✓	✓	✓
Backup and Recovery Suite						
Hitachi Multiplatform Backup/Restore			✓	✓		✓
TANTIA Technologies® HARBOR® Backup/Recovery and HYPERtape			✓	✓		✓
Nondisruptive database backup			✓	✓	✓	✓
Nondisruptive application backup			✓	✓	✓	✓
Data Movement Suite						
Hitachi RapidXchange™			✓	✓		✓
Hitachi Data Migration Services		✓	✓	✓		✓
Hitachi e-Copy			✓	✓	✓	
TANTIA Technologies HARBOR File Transfer			✓	✓		✓

The Business Value of Hitachi Freedom Storage Software Solutions

The HiCommand Management Framework delivers the ability to achieve centralized, worldwide control of the storage environment. Customers have the freedom to choose the precise solution—or combination of solutions—appropriate for their environment in four important areas, with each area targeted to have specific business value. In keeping with the Hitachi commitment to open solutions, Hitachi Freedom Storage Software Solutions include not only software solutions provided by Hitachi Data Systems, but also solutions provided by software vendors.

Storage Area Management Solutions – Hitachi Freedom Storage software and hardware solutions are managed through the powerful HiCommand Management Framework. This management tool substantially lowers TCO by enabling storage hyper-consolidation AND the intelligent remote management of critical information. The business value of centralized, global storage management is to increase business agility through heightened operational excellence and greatly reduced TCO. The HiCommand Management Framework is illustrated in Figure 1 and its ability to drive down TCO through centralized, worldwide control is illustrated in Figure 2. At the heart of Hitachi Freedom Storage Software Solutions are the Hitachi Resource Manager package and Performance Enhancement Solutions. Resource Manager simplifies IT operations, improves information availability, and enables truly centralized, worldwide control of critical enterprise information.





Figure 1 – HiCommand Management Framework allows management of multi-vendor hardware and software from the customer's platform of choice.

Performance Enhancement Solutions – The Performance Enhancement Suite increases performance service levels for business-critical OLTP and DSS applications, thereby enabling more rapid deployment of new applications by managing data more intelligently. For the Lightning 9900 V Series, both Resource Manager and the Performance Enhancement Suite have been enhanced to support the new virtual Fibre Channel Port technology and Host Storage Domains with the powerful Hitachi SANTinel™ and Priority Access programs.

Business Continuity Solutions – Hitachi Data Systems has an industry-unique commitment to the paradigm of Continuous Business. Continuous Business refers to an enterprise's ability to minimize system downtime, whether it is planned or unplanned. This is distinct from Business Continuity, which refers to the ability to recover from unplanned downtime. Building on the Hitachi Data Systems reputation for bulletproof reliability, the Lightning 9900 V Series, Thunder 9500 V Series, and Thunder 9200 storage systems offer complete redundancy and hot-replaceable components, delivering maximum uptime. To these already-robust platforms, Hitachi Data Systems adds Business Continuity solutions to ensure quick recovery from unplanned downtime resulting from acts of nature, human errors, application errors, and malicious attacks. Hitachi Data Systems is the only storage vendor that is providing nondisruptive copy solutions that ensure fast recovery and full data integrity. For the Lightning 9900 V Series, TrueCopy software performance has been improved to provide track instead of cylinder transfers and *resynch* time has been cut in half. For Lightning 9900 V Series systems, ShadowImage copy performance has improved with an increase of more than 40 percent in copy transfer rate and twice the number of pairs (4096).

Figure 2 – HiCommand Management Framework allows global systems management of Hitachi Freedom Storage hardware and software solutions.



Backup and Recovery Solutions – Backup/Restore is the largest cost element in storage ownership and the most important to high-availability computing. In today's global IT environment, it is no longer enough to have a backup copy of one's data; businesses must be able to perform backup and restore in the shortest possible time and with minimum disruption to information availability. Hitachi Data Systems therefore combines its own leading-edge Business Continuity software with those of industry-leading software partners to address customers' backup and restore requirements with solutions that reduce TCO. Hitachi Multiplatform Backup/Restore delivers high-performance, volume-level backup of open systems data. In addition to the HARBOR Backup/Recovery and HYPERTape Backup/Restore software packages, Hitachi Data Systems offers several features and services to achieve this goal. For example, ShadowImage and TrueCopy software for S/390 environments can be used to create a snapshot copy for near-instantaneous backup.

All Hitachi Freedom Storage backup and restore solutions are designed to meet the following customer requirements:

- Maximum reliability
- Centralized administration
- File-level backup and restore
- Online backup of popular databases, applications, and messaging services
- Out-of-the-box agents for popular servers such as Oracle® Server and BMC Software™
- Multiplatform backup to a mainframe or open systems server
- High-performance movement of data
- Ability to "snapshot" the backup source in UNIX®, Microsoft® Windows NT®/Windows® 2000 or IBM MVS®
- Ability to simulate or test the recovery scenario.

Data Movement Solutions – Sharing and moving large amounts of data is the key to business success in today's global information processing environments, which require that corporate data be readily available to all worldwide users, regardless of their location or the nature of their computing environment. With core business systems residing on a mix of mainframe and open systems, sharing and moving data among multiple, heterogeneous systems is key to fulfilling this requirement. Data from heterogeneous systems needs to be exchanged regularly, either for data synchronization between discrete applications or, increasingly, for populating data warehousing systems.

RapidXchange software offers high-performance and high-reliability data sharing between heterogeneous host platforms, such as IBM S/390®, UNIX, or Windows NT or Windows 2000. With RapidXchange software, high-speed data transfer is achieved without putting additional data load on the network infrastructure or tape transport equipment. Working with RapidXchange software, HARBOR File Transfer adds automation to the process of transferring large data files, at high-channel speeds, in either direction between open systems and mainframe servers. It can be thought of as an ultra-fast FTP transfer, except data is transferred over an IBM ESCON® instead of Ethernet link. HARBOR File Transfer enables UNIX and Windows applications, such as enterprise resource planning and decision support systems, to work with up-to-the-minute data extracted from S/390 environments. The Hitachi Data Migration Service helps migrate data from existing systems to newly installed systems while minimizing the impact on mission-critical applications. For the Lightning 9900 V Series, the Hitachi e-Copy feature provides server-free transfer of data between devices whether it is disk to disk or disk to tape.

Other Information Sources Available from Hitachi Data Systems

Two other publications that discuss the core competencies upon which the TrueNorth vision and strategic direction are based are available from Hitachi Data Systems. *The Lightning 9900 V Series Architecture Guide*, the *Thunder 9200 Architecture Guide*, and the *Thunder 9500 V Architecture Guide* describe these leading-edge products' key features and business values relating to respective product classes. These documents are downloadable in PDF format from <http://www.hds.com>.

If further information is required on any of the topics covered in these three guides, Hitachi Data Systems representatives can provide specialized presentations, reports, and expert knowledge.

Storage Area Management and Performance Enhancement Solutions

2

HiCommand™: Hitachi's Framework for Open Storage Management

Storage management is a complex problem that must support a myriad of existing components and new technologies. With the advent of storage area networks (SANs) and network attached storage (NAS), customers have greater flexibility and control over their storage environments than they have ever had in the past. However, this has also dramatically increased their storage management burden. These new, complex storage topologies lack a clear and consolidated approach to management. It is no longer good enough to provide basic services such as LUN management, device monitoring, and backup; now, it is critical to provide integrated management of all aspects of storage across a heterogeneous environment made up of multiple devices, such as SAN switches, complex storage systems, and storage appliances.

Customers have also invested heavily in "point" solutions and systems management applications for their current environments. It is not sufficient for a vendor to provide a replacement for existing technologies; the vendor must provide a complete solution by allowing customers to leverage their existing investments, whether they are in hardware, software, or people.

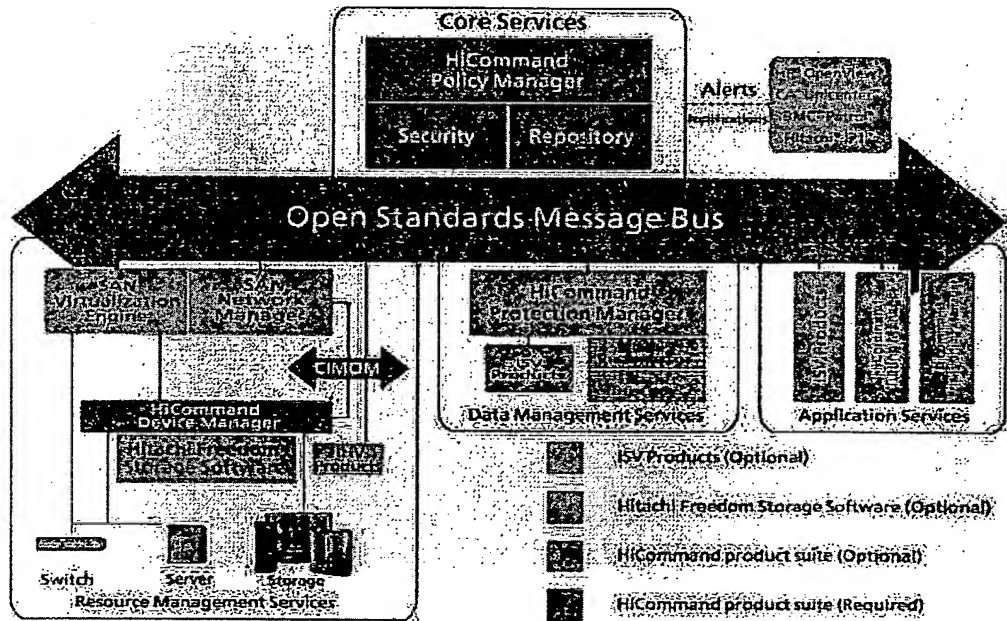
No Single Vendor Can Meet All Storage Requirements for Today's Enterprises

Hitachi provides an open platform, based on industry-standard management protocols, along with high, value-add functional components, such as policy-based automation and heterogeneous device management, allowing customers to incorporate multiple vendor offerings to create a truly complete solution for managing and monitoring their storage environments.

As shown in Figure 1, the HiCommand Management Framework supports a modular, building-block approach to designing and addressing business-specific storage management requirements. As a result, the HiCommand Management Framework enables customers to build the most suitable management environment based on a combination of Hitachi products and best-of-breed storage hardware and software. The HiCommand Management Framework is customizable to integrate easily with existing infrastructures and accommodate future growth plans.

The HiCommand Management Framework's policy-based automation module (Figure 3) simplifies management functions and enables integrated administration of multi-vendor storage environments. Furthermore, standard storage management operations come pre-loaded in the policy-based automation module. The HiCommand Management Framework also provides definition tools and an easy-to-use graphical user interface (GUI) interface for automating more complicated storage management processes, resulting in lower TCO.

Figure 3 – The HiCommand Management Framework is based on CIM and SOAP industry standards.



Open, Standards-based Approach

At the foundation of HiCommand Management Framework is a message bus that allows individual functional components to "plug and play." Functional components are incorporated into the overall solution based on a usage of the Common Information Model (CIM), and the Simple Object Access Protocol (SOAP) provides a standard envelope for a messaging bus that enables simple plug-in of functional modules. This model provides a clear set of information content for management and monitoring of the entire storage and system environment. Using this object-model-based approach, solutions can readily communicate with all network-attached devices and perform a variety of configuration, administration, and other functions, allowing for a complete, integrated hardware and software management environment. Adopting this model is the key to overall interoperability in the storage and systems management environments. If a solution component adheres to this model, then its management or incorporation will be relatively straightforward.

Another industry standard, eXtensible Markup Language (XML), is fast becoming the norm for data communication within an integrated environment. Using XML to express the messages within the CIM model enables not only a common understanding of the required content, but also a clear and universal understanding of the message format and flexible application integration. Hitachi is creating an open and extensible architecture using these two foundational standards.

Hitachi is creating an open and extensible architecture using these three foundational standards. A more thorough discussion of the HiCommand Management Framework initiative and the open industry standards can be found at http://www.hds.com/pdf/PERI116-00_WhtPpr.pdf

XML is fast becoming the norm for data communication within an integrated environment.

HiCommand Allows Global Systems Management of Multi-vendor Storage Systems

At the heart of the Hitachi Storage Area Management Solutions are the HiCommand Management Framework, its two modules, HiCommand Device Manager and HiCommand Tuning Manager software, Hitachi Resource Manager™ software, and the Hitachi Performance Enhancement Solutions. Resource Manager software simplifies IT operations, improves information availability, and enables truly centralized, worldwide control of critical enterprise information. Hitachi Performance Enhancement Solutions increase performance service levels for business-critical OLTP and DSS applications, thereby enabling more rapid deployment of new applications by managing data more intelligently. For the Hitachi Freedom Storage™ Lightning 9900™ V Series, both Resource Manager software and Hitachi Performance Enhancement Solutions have been enhanced to support the new virtual Fibre-channel port technology and Host Storage Domains in combination with the powerful Hitachi SANTinel™ software and Hitachi Priority Access programs.

Hitachi Freedom Storage software and hardware solutions are administered through the powerful HiCommand™ Management Framework, which substantially lowers TCO by enabling storage massive consolidation AND the intelligent remote management of critical information. The business value of centralized global storage management is to increase business agility through increased operational excellence and greatly reduced TCO.

Hitachi Data Systems has long recognized the importance of independent software vendors (ISVs) and independent hardware vendors (IHVs), and is working with vendors whose offerings are complementary. This is accomplished through the HiCommand open APIs and through joint development and engineering, joint certification, and where appropriate, marketing agreements. At this time, there are over 75 vendors supporting the HiCommand open API. Both Hitachi Data Systems and its software strategic alliance members believe this open architecture will best benefit customers by allowing them to manage their storage through their vendor of choice.


The Components of HiCommand Device Manager Software

HiCommand Device Manager software consists of three components:

1. The **HiCommand server** resides on any Microsoft® Windows NT®/Windows® 2000, or Sun™ Solaris™ server, and connects to Lightning 9900 V Series, Hitachi Freedom Storage Thunder 9500™ V Series, or Hitachi Freedom Storage Thunder 9200™ storage system(s) over a private LAN. A single HiCommand server can manage multiple Hitachi Freedom Storage systems, as well as Sun StorEdge™ 9900 Series and T3 storage systems, providing a common point of control that delivers:
 - Out-of-band SANs, NAS, or direct attached storage (DAS) connections that enable data to run between the server and storage at lightning-fast speeds
 - Seamless integration between Device Manager software and solutions from industry leaders such as Sun Microsystems™, BMC Software™, IBM®, and Computer Associates® through use of XML-based APIs
 - Remote and local access control that offers a centralized approach to managing storage from any location
2. A **GUI, browser-based interface** is easy to learn for simple and efficient storage administration. The HiCommand browser (Figure 4) is intuitive and easy to use. A browser from anywhere can control the HiCommand server. In addition, Hitachi Data Systems also supplies an optional Hitachi Command Control Interface (CCI), which enables users to perform Hitachi TrueCopy™ and Hitachi ShadowImage™ operations by issuing commands from the UNIX®/PC server host to the Hitachi Freedom Storage system. The CCI software interfaces with the

HiCommand enables an enterprise to manage multi-vendor systems and software solutions.

Figure 4 – The HiCommand GUI browser is intuitive and easy to use.



Out-of-band storage area networks (SANs) enable data to run between the server and storage at lightning speeds.

system software and high-availability software on the UNIX/PC server host as well as the TrueCopy and ShadowImage software on the Lightning 9900 V Series systems and the Thunder 9570V™ system. The CCI software provides failover and other functions such as backup commands to allow mutual hot standby in cooperation with the failover product on the UNIX/PC server. CCI also supports a scripting function that allows users to define multiple TrueCopy and/or ShadowImage operations in a script (text) file. Using CCI scripting, administrators can set up and execute a large number of TrueCopy and/or ShadowImage commands in a short period of time while integrating host-based, high-availability control over remote copy operations.

3. **HiCommand host agents** reside on the customer's application servers and "push" information back to HiCommand Device Manager, including percentage of file utilization of LUNs, host bus adaptors (HBAs), World-Wide Names (WWNs), operating system SCSI addresses, and other useful device-specific and file-level information.

How Hitachi Resource Manager Supports Business Objectives

To be successful in today's business environment, companies must attain the highest levels of business agility, with maximum IT system efficiency and performance. Today, there are three major storage architectures—DAS, NAS, and SAN—and many types of servers. The "storage-centric" model of computing is widely accepted today. In this model, enterprise data is the most important element of the IT infrastructure. Intelligent storage and network management software suites like Resource Manager, therefore, are at the heart

of the enterprise's ability to achieve high levels of performance and availability to meet business objectives. Ever-expanding storage in both centralized and distributed environments must be managed intelligently while ensuring that storage data remains available to all who need it.

Hitachi Resource Manager Is Modularly Designed

The complexities of heterogeneous computing environments make storage management a daunting task. Hitachi Data Systems offers the ideal solution to this difficult challenge – the Resource Manager package. As shown in Table 2, this comprehensive systems and software management offering brings together 11 powerful software solutions for either the open systems or IBM S/390®-compatible suites.

	Product Support				OS Support	
	Thunder 9531V™ Thunder 9532V™ Thunder 9533V™	Thunder 9570V™	Lightning 9900™ Series	Lightning 9900 V Series	Open Systems	IBM S/390®
Hitachi Resource Manager						
Remote Console – Storage Navigator	✓	✓	✓	✓	✓	✓
Hitachi Graph-Track™			✓	✓	✓	✓
SNMP Agent	✓	✓	✓	✓	✓	✓
Sun™ Java™ Agent	✓	✓	✓	✓	✓	✓
Host Storage Domains	✓	✓		✓	✓	
LUN Manager / LUN Expansion	✓	✓	✓	✓	✓	
Hitachi SAnTinel™	✓	✓	✓	✓	✓	✓
Virtual LUN / LUN Expansion	✓	✓		✓	✓	
Virtual Logical Volume Image Manager (VLVI)			✓	✓		✓
Hitachi FlashAccess™			✓	✓	✓	✓
Cache Management Host Agent			✓	✓		✓

Table 2 – Summary of Hitachi Resource Manager solutions and the products they support.

Shipped with every Thunder 9200, Thunder 9500 V Series, or Lightning 9900 V Series storage systems, these integrated software packages are used by storage managers to display system configuration, create user name/password security for administrators, set up RAID groups, allocate LUNs, expand LUNs, and format storage. Storage administrators who need a flexible, easy-to-use configuration and error management tool will find the Resource Manager package an ideal choice. Users can monitor and manage their storage systems through a GUI.

The Resource Manager package can also help optimize performance by providing valuable resource utilization information, such as I/O activity, cache usage, and availability status/event notification. Administrators can use password protection to authorize users. This provides protection against unauthorized access to the management mode of Resource Manager software, and stores the password list at the array. Reliable error management and reporting offers a big breakthrough for IT troubleshooters who seek rapid-fire detection and notification to maintain productivity.

The following is a summary list of the many ways in which the Resource Manager package can simplify IT operations, improve availability and performance service levels for business critical OLTP and DSS applications, and thereby enable the more rapid deployment of new applications by managing data more intelligently.

Service-level Management

- Policy management
- Security management
- Automation
- Monitoring/reporting
- Storage management

Hitachi Resource Manager provides the ultimate in reliable, easy-to-use hardware performance and availability management software.

Deployment Management

- Asset management
- Change management
- Capacity planning

Compliance Management

- Event management
- Performance management
- Quota planning
- Accounting management
- Problem planning

1. Remote Console – Storage Navigator

The Remote Console for Lightning 9900 V Series systems – Storage Navigator is provided as a Java applet program, which can execute on any machine that supports a Java Virtual Machine (JVM). The Remote Console – Storage Navigator PC hosts the Storage Navigator Java applet program and is attached to the Lightning 9900 V Series system(s) via a TCP/IP LAN. When a Storage Navigator accesses and logs into the desired service processors (SVPs), the Remote Console applet is downloaded from the SVP to the Remote Console and runs on the Web browser of the Remote Console PC. In this way, the Remote Console communicates with the attached Lightning 9900 V Series or Thunder 9500 V Series systems via a TCP/IP network.

Two LANs can be attached to the Lightning 9900 V Series or Thunder 9500 V Series storage systems: an internal LAN (private LAN), which is used to connect the SVPs of multiple storage systems, and the user's intranet (public LAN), which allows access to the Storage Navigator functions from individual Storage Navigator PCs. The Remote Console – Storage Navigator communicates directly with the SVP of each attached storage system to obtain storage system configuration and status information and send user-requested commands to the storage system. The Lightning 9900 V Series storage system's Storage Navigator Java applet program is downloaded to the Storage Navigator (Web client) from the SVP (Web server) each time the Storage Navigator is connected to the SVP. The Storage Navigator Java applet program runs on Web browsers, such as Internet Explorer® and Netscape Navigator®, which run under the Windows 2000, Windows NT, and Solaris operating systems to provide a user-friendly interface for the Lightning 9900 V Series system's Storage Navigator functions.

2. Graph-Track Software Takes the Guesswork Out of Performance Management

As the cornerstone of the software utilities of the Resource Manager software suite, the Graph-Track GUI is widely recognized as user friendly. Icons and pull-down menus define and display data, while point-and-click capabilities and online help further enhance its ease of use. All Graph-Track data can be exported in several formats quickly and effortlessly, for use in other data analysis and reporting programs.

Graph-Track software gives users a more reliable and centralized way to manage performance. Available for S/390, UNIX, and Windows NT/2000 environments, this unique tool monitors hardware performance and supplies complete system storage information at the touch of a button. The robust

The Hitachi Graph-Track GUI is widely recognized as user friendly.

GUI-oriented command screen for Graph-Track software is shown in Figure 5. Running on a PC attached to a dedicated LAN, Graph-Track software displays real-time or historical data for all connected Lightning 9900 V Series systems, helping users identify important peaks in utilization and ongoing trends in processing. With Graph-Track software, users can analyze discrete storage activity and determine the precise impact of each activity on system operations. It scrutinizes activity all the way to the logical-device level, reporting on channel interface processor usage rates, cache usage rate by function, and logical-device utilization. This data pinpoints specific activities that may have a negative impact on operations and lets the user fine-tune the system so that it can be used to maximum advantage.

Graph-Track software also highlights valuable cache information. Cache read/write and read-hit ratios are reported in real time. Each system connected to Graph-Track software can be tuned separately.




Figure 5 – The GUI interface on Hitachi Graph-Track software simplifies performance management.

3. SNMP Support Ensures Compliance with All Major Systems Management Platforms

The Resource Manager package supports the simple network management protocol (SNMP). SNMP is the most widely used network management protocol in the marketplace. System information messages (SIMs) are converted to SNMP management information blocks (MIBs) for use with Resource Manager software, or the VERITAS® Storage Manager™, CA Unicenter® or HP® OpenView®. With CA Unicenter, “super MIBs” allows the control of the Lightning 9900 V Series, Thunder 9500 V Series, or Thunder 9200 storage systems using CA Unicenter commands. This out-of-band¹ reporting format ensures that Resource Manager software can determine the status of your Hitachi Freedom Storage system even if the data channel path is down.

¹ “Out-of -band” refers to the path that system management messages travel, which is a different path from the data. “In-band” messages would compete with data for bandwidth.


4. Sun™ Java™ Agent

The Resource Manager Java agent allows the Remote Console – Storage Navigator Java applet program to run on Web browsers, such as Internet Explorer and Netscape Navigator. These programs run under Windows 2000, Windows NT, and Solaris operating systems and provide a user-friendly interface for the Lightning 9900 V Series and the Remote Console – Storage Navigator functions for Thunder 9500 V Series, or Thunder 9200 storage systems.

5. Host Storage Domains

The Lightning 9900 V Series and Thunder 9500 V Series systems both provide for “virtual” Fibre Channel ports that are logically managed by intelligent Fibre Channel controller cards. Each physical Fibre Channel port has multiple Host Storage Domains (HSDs). An HSD supports a “logical” Fibre Channel port (Fig. 6), each with its own set of LUNs. Hosts are matched to their assigned HSD based on a unique WWN identifier. LUN security is provided by SANTinel software, which fences access to LUNs based upon host WWNs. Under this innovative software structure, one LUN is reserved for the command and control device.

Figure 6 – Host Storage Domains reduce costs through fewer physical connections.



6. Hitachi LUN Manager and LUN Expansion Manager Software Simplify Configuration Management and Reduce Staffing Costs

LUN Manager is an open systems management utility. With LUN Manager software, open systems LUNs can be defined, configured, and maintained. There is no more waiting for the hardware vendor to come and make configuration changes. LUN Manager includes an easy-to-use, GUI-based interface that allows the definition of paths for LUNs, the reconfiguration of LUN-to-port assignments, or the viewing of the Lightning 9900 V Series or Thunder 9500 V Series storage systems' remote service information messages. Because LUN Manager software can assign multiple paths to a single LUN, support of alternate path failover, path load balancing, and clustered systems is possible. Running on a standard Windows-based PC connected to the storage systems by a dedicated LAN, LUN Manager can support up to eight Lightning 9900 V Series or Thunder 9500 V Series storage systems.

Hitachi LUN Expansion Manager software features logical unit size expansion that dramatically improves LUN flexibility. Up to 36 physical LUNs can be presented to the operating system as a single, large LUN. This capability simplifies storage management because there are fewer LUNs to manage. Two levels of password protection, user and administrator, ensure maximum data security. Users can view only Lightning 9900 V Series or Thunder 9500 V Series system LUN configuration information, while administrators can access all LUN configuration information and functions. Administrators can customize access privileges for individual users, thus providing maximum flexibility and control of these powerful management capabilities.

With Hitachi LUN Manager software, support of alternate path failover, path load balancing, and clustered systems is enhanced.

7. SANtinel Software

SANtinel software controls host access to Hitachi Freedom Storage LUNs in open systems, multi-platform, or SAN environments. This enables users to restrict server access to only those LUNs for which they are authorized, as shown in Figure 7.

SANtinel software for open systems, running on the Lightning 9900 V Series or Thunder 9500 V Series systems, allows administrators to define multiple Host Storage Domains for each physical Fibre Channel port on the system. A Host Storage Domain consists of a virtual (logical) Fibre Channel port and an associated set of LUNs visible only via that logical Fibre Channel port.

Connections from host servers arriving at the physical Fibre Channel port are routed to the logical Fibre Channel port within the appropriate Host Storage Domain based upon the Fibre Channel WWN of the host.

The storage administrator creates the Host Storage Domain, sets the host mode of the logical Fibre Channel port within that Host Storage Domain, and maps the LUNs to appear within that Host Storage Domain. Then the storage administrator authorizes a particular host to "see" a Host Storage Domain on that physical port by assigning the Fibre Channel WWN of that host to the Host Storage Domain in question.

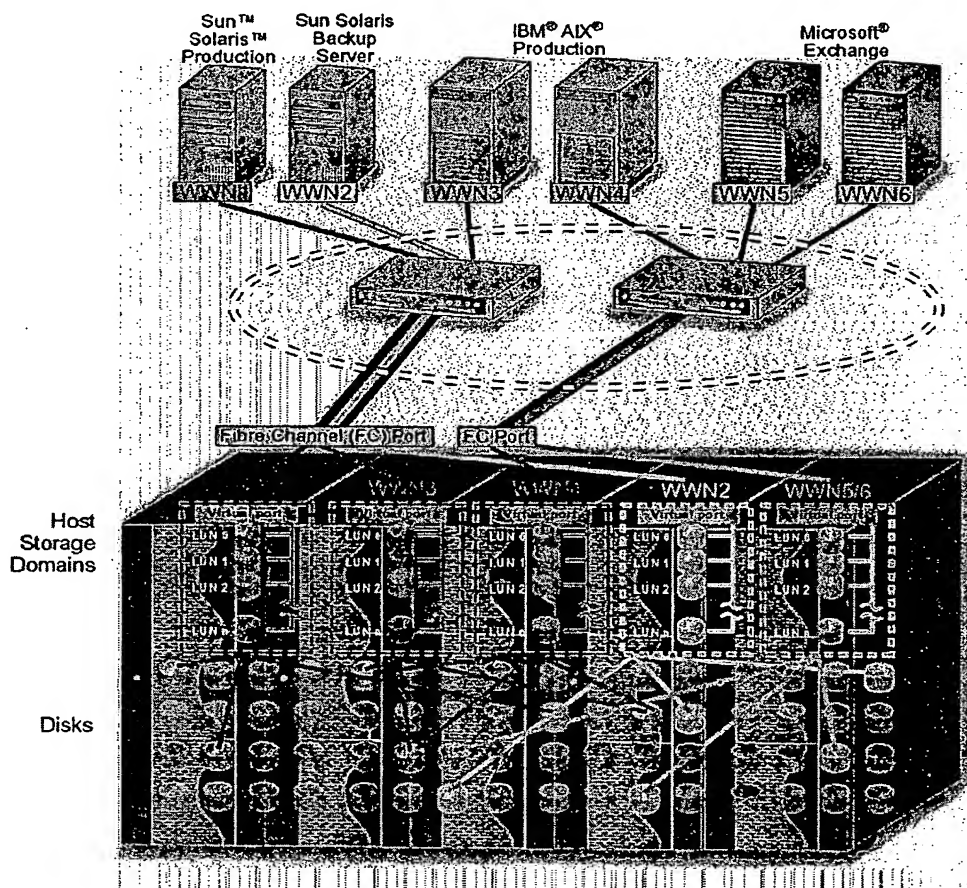
Because the WWN is used to route connections to the appropriate Host Storage Domain, the WWN may only be assigned to a single Host Storage Domain on each Fibre Channel port.

SANtinel software may be used to authorize multiple hosts to access a particular Host Storage Domain. SANtinel software may be further used to permit or deny a particular host to access individual LUNs within that Host Storage Domain, based upon the WWN of the host.

SANtinel software enables users to restrict server access to only those LUNs for which they are authorized.

Hitachi Virtual Logical Volume Image Manager improves performance.

Figure 7 – SANTinel software uses pre-defined zones to protect data from unauthorized access.



SANTinel software may be used to create up to 128 Host Storage Domains per physical Fibre Channel port. Each Host Storage Domain may contain up to 256 LUNs, with the maximum number of LUNs contained within all Host Storage Domains on a given Fibre Channel port reaching 512.

In a similar fashion to open systems, SANTinel software for S/390 environments allows administrators to restrict S/390 host access to the logical devices (LDEVs) on Lightning 9900 V Series systems. Each LDEV can be set to communicate only with user-selected host(s). SANTinel software for S/390 prevents other hosts from seeing the secured LDEV and from accessing the data contained on the secured LDEV. The licensed SANTinel software for S/390 Storage Navigator displays SANTinel information and allows administrators to perform SANTinel operations.

8. Virtual LUN and LUN Expansion Software

Virtual LUN software allows administrators to convert fixed-size volumes into several smaller, variable custom-sized volumes. Using the Remote Console – Storage Navigator software, users can configure custom-size volumes by assigning a logical address and a specific number of cylinders/tracks (for S/390 data) or MB (for open-systems data) to each custom logical volume image (LVI)/LUN. Virtual LUN software improves data access performance by reducing logical-device contention as well as host I/O queue times, which can occur when several frequently accessed files are located on a single volume. Multiple LUN types can

be configured within each array group. Virtual LUN enables administrators to more fully utilize the physical storage capacity of the Lightning 9900 V Series or Thunder 9500 V Series storage systems, while reducing the amount of administrative effort required to balance I/O workloads.

When Virtual LVI/LUN software is used in conjunction with FlashAccess software, users can achieve even better data access performance than when either Virtual LVI/LUN or FlashAccess software are used alone.

The LUN Expansion Manager allows administrators to create virtual LUNs that are larger than standard OPEN LUNs, by expanding the size of a selected LUN up to 36 times its normal size. The maximum size depends on the type of configuration. For example, an administrator can expand an OPEN-9 LUN to a maximum size of 265GB (7.3GB x 36). This capability enables open-systems hosts to access data on the entire Lightning 9900 V Series or Thunder 9500 V Series storage systems, using fewer logical units. LUN Expansion Manager allows host operating systems that have restrictions on the number of LUNs per interface to access larger amounts of data.

***Hitachi FlashAccess
software dramatically
improves performance.***

9. Virtual Logical Volume Image (VLVI) Manager for S/390 Environments Allows Very Large to Very Small Virtual Volume Configuration

Virtual Logical Volume Image (VLVI) Manager is a mainframe software utility that optimizes Lightning 9900 V Series system capacity utilization by allowing users to configure multiple *virtual LVIs* in place of an LVI. Data volumes as small as a single cylinder or as large as a full 3390-9 can be defined. Volume size is determined in cylinder increments. Each virtual LVI requires one physical address, with a maximum of 8192 addresses per Lightning 9900 V Series system. Different types of LVIs can coexist within an array group with no need to convert them to a common LVI. This feature maximizes array group capacity and enhances configuration flexibility. VLVI Manager improves performance by reducing logical device contention and operating system queuing. It also boosts remote copy performance by avoiding the need to copy the entire volume.

10. Hitachi FlashAccess Software Allows Data to Be “Locked and Unlocked” in Cache On-the-Fly

The FlashAccess software utility allows users to dynamically “lock” and “unlock” data into cache in real time. Read and write functions are then performed at cache speeds, with no disk latency delay. With FlashAccess software, a portion of cache memory can be allocated to specific data. Users can add, delete, or change FlashAccess managed data at any time, quickly and easily.

Defined by the LVI for mainframes or LUN for open systems, cache data can be as small as a single track or as large as an entire volume. For increased configuration flexibility, FlashAccess software offers multiple modes of operation. It can be used in conjunction with Hitachi RapidXchange™ software to increase the speed of data transfer and, therefore, improve performance of mainframe to open systems data exchange. RapidXchange software offers high-speed data transfer among IBM S/390, UNIX, Windows NT, and Windows 2000 platforms. RapidXchange software transfers can be open-to-open as well as mainframe-to-open.

11. Cache Manager Host Agent

The Cache Manager Host Agent enables mainframe users to perform FlashAccess operations on S/390 LVI's by issuing commands from the S/390 host system to the Lightning 9900 V Series system.

Table 3 – Summary of Hitachi Performance Enhancement Software Solutions and the products they support.

	Product Support				O/S Support	
	Thunder		Lightning			
	9200™	9500™ V Series	9900™ Series	9900 V Series	Open Systems	IBM® S/390®
Hitachi Performance Enhancement Solutions—Hitachi Performance Maximizer Package						
Hitachi Performance Monitor			✓	✓	✓	✓
Hitachi CruiseControl™			✓	✓	✓	✓
Hitachi Priority Access				✓	✓	
TPF/MPLF			✓	✓		✓

Hitachi Performance Enhancement Solutions

Hitachi Performance Enhancement Solutions increase performance service levels for business-critical OLTP and DSS applications, thereby enabling more rapid deployment of new applications by managing data more intelligently. For Lightning 9900 V Series systems, both the Resource Manager package and the Performance Enhancement Software Suite have been enhanced to support the new virtual Fibre Channel port technology and Host Storage Domains with the powerful SANtinel and Priority Access software programs.

Hitachi Performance Maximizer Package

Hitachi Performance Maximizer, a comprehensive package comprising Performance Monitor, CruiseControl, and Priority Access software, provides automated service-level management. The Performance Monitor console, a key component of the Performance Maximizer, interfaces with the remaining products to facilitate centralized storage management. Through a link-and-launch capability with CruiseControl and Priority Access software, the Performance Monitor tracks and displays detailed usage and trends, and enables automation through established policies and thresholds.

1. Performance Monitor Software

Performance Monitor is a powerful software utility that provides a variety of performance measurements for system administrators to optimize performance of their Hitachi Freedom Storage systems.

2. CruiseControl Software

For open systems volumes, CruiseControl software provides automatic performance tuning for Lightning 9900 V Series systems by automatically eliminating performance bottlenecks. Utilizing the many high-speed internal paths of Lightning 9900 V Series systems to optimize data placement, CruiseControl software automatically monitors, analyzes, and moves logical volumes to eliminate "hot spots" and provide load balancing to maintain predetermined performance levels. Independent of the type of host attachment, CruiseControl software makes recommendations for administrator approval in either automatic or assisted mode. CruiseControl software replaces time-consuming and sometimes error-prone, manual load balancing with simple automated procedures. This ensures long-range optimal performance and reduced cost of ownership.

CruiseControl software, when combined with FlashAccess software for

Hitachi CruiseControl automatically monitors, analyzes and moves logical volumes to eliminate "hot spots."

mainframes and open systems, allows customers to manage the cost of storage, as data is consolidated onto fewer, better-managed systems. Storage administrators can “fill it up and max it out,” confident that they can meet performance requirements all the way to 37TB raw capacity on a Lightning 9900 V Series system using 73GB drives. CruiseControl software then takes over and eliminates “hot spots,” thereby automating performance to service-level objectives, while FlashAccess software locks data into cache, guaranteeing data access at memory, not disk, speed.




Figure 8 – Priority Access software allows administrators to allocate bandwidth in line with quality of service.

With the massive consolidation made possible by the Lightning 9980V™ storage system's capabilities, customers will be able to not only lower the original purchase costs, but also get additional savings from lower storage management costs, lower software and maintenance costs, and reduced real estate and environmental costs. This is often the major payback for Hitachi Freedom Storage products. Reduced per-megabyte administration costs alone can provide a substantial return on investment.

3. Priority Access Software

Priority Access software allows open-systems users to designate prioritized ports (e.g., for production servers) and non-prioritized ports (e.g., for development servers) and set thresholds and upper limits for the I/O activity of these ports. This enables users to tune the performance of the development server without affecting the production server's performance as illustrated in Figure 8.

Transaction Processing Facility/Multi-path Lock Facility (TPF/MPLF)

The Lightning 9900 V Series systems support the Multi-path Lock Facility (MPLF) for the IBM highest performance transaction processing software environment – the Transaction Processing Facility (TPF). TPF is used by many of the world's largest customer reservation systems and by the world's largest financial institutions. In either native TPF mode or under VM, MPLF provides extremely high performance record-level locking, which enables high levels of concurrent data access across multiple channel paths. RAID-1, RAID-5, 3390-3, and 3390-9 LVIs are supported.

Hitachi Parallel Access Volume (HPAV) and Multiple Allegiance (MA)

Hitachi Parallel Access Volume (HPAV) software enables the S/390 host system to issue multiple I/O requests in parallel to single logical devices (LDEVs) in a Lightning 9900 V Series system. HPAV software can provide substantially faster host access to the S/390 data stored in a Lightning 9900 V Series system. The Workload Manager (WLM) host software function enables the S/390 host to utilize the HPAV functionality of a Lightning 9900 V Series system. The Lightning 9900 V Series systems support both static and dynamic HPAV functionality. Multiple Allegiance (MA) software extends this capability to applications running on multiple S/390 servers. Together, HPAV and MA software reduce queuing, which results in significantly decreased batch times and lightning-fast responses in today's high-transaction environments.

The VERITAS Foundation Suite helps decrease the total cost of operations.

VERITAS Software Simplifies Management and Increases Availability

The VERITAS Foundation Suite™ combines the VERITAS Volume Manager™ and VERITAS File System™ to simplify storage management while ensuring high availability for critical systems. VERITAS Volume Manager is an easy-to-use, online disk and storage management tool for enterprise computing environments. VERITAS File System is an enterprise-class, journaling file system that facilitates high performance, quick recovery, and easy scalability. The VERITAS Foundation Suite provides a strong foundation, not only for managing current storage requirements, but also (via its open architecture) for leveraging future technologies and opportunities. Downtime, whether scheduled or unscheduled, is costly. The VERITAS Foundation Suite helps decrease the total cost of operations and brings greater stability to vital information systems.

VERITAS Volume Manager increases availability and lowers costs.

1. VERITAS Volume Manager

VERITAS Volume Manager increases availability by creating redundant storage solutions for critical data. At the same time, it helps lower storage costs by making the most of the storage already in-house. "Hot relocation" capabilities automatically construct a mirror on free space in the event of device failure. This capability is ideal for those applications that require a non-hardware RAID solution (that is, software-based RAID-0 or RAID-5). Hot spare devices are identified for use in case there is a disk failure, but if they are unavailable, any available free space is used to maintain redundancy. If the server hardware and operating system support hot swapping, data can be swapped to the new device while the system is running.

Instead of adding to overall storage costs, the VERITAS Foundation Suite manages storage more efficiently for overall cost savings and improves the performance of existing storage. By providing the flexibility and scalability to add and manage new storage easily, it allows users to take a proactive approach to storage management. VERITAS Volume Manager enables the effective use of current hardware by identifying available segments or partitions on storage devices and combining them into one or more virtual disks or logical volumes online, without interruption.

This feature also allows the striping of data across multiple physical volumes. Striping reduces the time spent by the application waiting for disk head movements and generates faster overall system performance. In addition, data can be mirrored between multiple RAID devices. Online load balancing between RAID logical units, along with the reconfiguration of RAID logical units, can also be performed without interrupting access to data.

VERITAS Volume Manager performs regular maintenance tasks online while maintaining system and data availability. These tasks include data backup and defragmentation, storage allocation resizing, storage additions, faulty or slow device replacement, and device reconfiguration for I/O load balancing or new application support. Availability is further maintained via the Dynamic MultiPathing (DMP) capabilities of VERITAS Volume Manager. If a path fails or becomes unavailable, VERITAS Volume Manager automatically reroutes data to the open path(s).

2. VERITAS File System

VERITAS File System works to reduce outages due to file system “panics.” If it detects a disk or file system error, VERITAS File System isolates the problem without causing a failure on the whole file system. If an outage occurs, VERITAS File System can restore maximum access to data in seconds. All writes to the file system are recorded in a journal or intent log until they are written to index or data blocks. During the restart after failure, VERITAS File System replays the log record and completes any uncompleted writes within seconds. Its online snapshot, supported by common backup packages, enables full backup without data lockout.

When time is money, it pays to ensure optimized application performance, especially for high I/O applications. Performance can be significantly increased with the features of VERITAS File System and VERITAS Volume Manager. By allocating storage in extents (groups of contiguous blocks) rather than in fixed block sizes, VERITAS File System reduces fragmentation and requires fewer indirect pointers to files. This reduces seek time due to disk latency. Reading/writing to larger extents rather than to smaller blocks requires fewer I/O operations. VERITAS File System also facilitates fast performance by automatically tuning I/O to match VERITAS Volume Manager data layout. VERITAS Volume Manager dramatically improves overall I/O throughput on disk controllers that have two separate I/O paths by using both paths simultaneously. By configuring the VERITAS File System to do this automatically, an immediate performance improvement in existing applications that are I/O intensive will result. If there are storage locations with very high access patterns, VERITAS Volume Manager will establish a “sparse mirror.” A “sparse mirror” is a mirror of a small part of the overall data. This reduces contention on read requests by making another copy available. If the sparse mirror resides on storage with very fast access characteristics, such as on the Lightning 9900 V Series and Thunder 9500 V Series systems, then the access provided is even faster. In addition, FlashAccess software for open systems and S/390 platforms for the Lightning 9900 V Series systems provides for dynamic cache residency, which can further enhance performance.

VERITAS File System reduces fragmentation and requires fewer indirect pointers to files, reducing seek time.

3. VERITAS Storage Administrator Improves System Manageability

When an enterprise has limited time and resources to monitor storage, centralization of data management can provide maximum benefits to operations efficiency and reduced cost efforts. The VERITAS Foundation Suite along with Hitachi Freedom Storage systems allow the centralization, automation, and integration of critical resources—whether local or distributed, online or offline. VERITAS Foundation Suite is easy to use, reducing the training costs and the time spent completing administrative tasks.

The VERITAS Foundation Suite and Hitachi Freedom Storage systems allow the centralization, automation, and integration of critical resources.

Simplified remote storage management is possible through VERITAS Storage Administrator, a Java-based GUI that ships with the VERITAS Foundation Suite. The GUI helps administrators defragment storage, identify and group devices, implement mirroring and striping, combine segments from multiple devices into logical volumes, and resize or move volumes on physical disks. All of this occurs while storage remains available and users remain online. Systems managers can monitor I/O traffic to identify bottlenecks and inefficiencies with the VERITAS Foundation Suite command line interface. VERITAS Volume Manager further contributes to manageability in a storage solution by making it easy to move data between nodes in a SAN environment. Finally, the VERITAS File System allows a storage administrator to group large collections of files for commercial applications, eliminate media-size limitations, and easily manage space requirements for a large number of user accounts.

4. VERITAS SANPoint Control™

VERITAS SANPoint Control is the data storage industry's most advanced, easy-to-use, centralized management tool for automatic discovery, visualization, and zoning administration of SAN-connected devices. Leveraging the VERITAS breakthrough V3™ technology and existing industry standards, VERITAS SANPoint Control provides the first drag-and-drop zoning capabilities in a heterogeneous SAN. Building on the unique storage visualization capabilities of VERITAS Volume Manager, VERITAS SANPoint Control enables a single point of management from LUN to interconnect and to SAN-connected hosts. Host-, device- or fabric-based maps provide customizable, intuitive views into your SAN as well as real-time events to diagnose and avoid outages. By masking the administrative complexity of the SAN environment, VERITAS SANPoint Control allows the potential of SANs to be realized today.

Business Continuity Solutions

3

Hitachi Data Systems has an industry-unique commitment to the paradigm of Continuous Business. Continuous Business refers to an enterprise's ability to minimize system downtime whether it is planned or unplanned. This is distinct from Business Continuity, which refers to the ability to recover from unplanned downtime. Building on the Hitachi Data Systems reputation for bulletproof reliability, the Hitachi Freedom Storage™ Lightning 9900™ V Series, Thunder 9500™ V Series, and Thunder 9200™ storage systems offer complete redundancy and hot-replaceable components, delivering maximum uptime. To these already-robust platforms, Hitachi Data Systems adds Business Continuity software solutions to ensure quick recovery from unplanned downtime resulting from acts of nature, human errors, application errors, and malicious attacks. Hitachi Data Systems is the only storage vendor that is providing non-disruptive copy solutions that ensure fast recovery and full data integrity. For the Lightning 9900 V Series and Thunder 9570V™ storage systems, Hitachi TrueCopy™ software performance has been improved to provide track instead of cylinder transfers and *reync* time has been cut in half. For Hitachi ShadowImage™ software, copy performance has improved with a copy transfer rate increase of more than 40 percent and twice the number of pairs (8192).

Increased IT Service Levels in Availability and Performance through Business Continuity Software Solutions and Services

In today's global IT environment, performance and availability go hand in hand. If a Web site is slow in delivering requested information, it might as well be offline. Of course, if a Web site is down for any length of time, the business may not be around for long. That's why the Lightning 9900 V Series, Thunder 9500 V Series, and Thunder 9200 storage systems offer complete redundancy and hot-replaceable components, delivering maximum uptime.

It is important to understand the terminology used in discussions of Business Continuity and Disaster Recovery since both are similar terms used in discussions of high-availability computing. Continuous Business refers to an enterprise's ability to minimize system downtime whether it is planned or unplanned. This is distinct from Business Continuity, which refers to the ability to recover from unplanned downtime. Business Continuity Planning software and services is an "umbrella" term covering both disaster recovery planning and business resumption from a business interruption.

A business interruption is any event, whether anticipated (i.e., public service strike) or unanticipated (i.e., blackout) that disrupts the normal course of business operations at an enterprise location. A disaster is any event that creates an inability for an organization to provide critical business functions for an undetermined period of time. While a disaster is a business interruption, not all business interruptions are disasters.

Disasters can be subdivided into three categories:

Man-made Disasters

- Power outages
- Fires
- Explosions
- Accidents
- Bombings
- Sabotage
- Building collapse
- Burst pipes
- Environmental contamination

The Lightning 9900 V Series and Thunder 9500 V Series storage systems deliver unsurpassed uptime and performance.

High availability terminology is important to understand.

Natural Disasters

- Power outages
- Fires
- Earthquakes
- Tornadoes and wind storms
- Volcanic eruptions
- Hurricanes, monsoons, and typhoons

Man-made Political Disasters

- Riots and civil disturbances
- Terrorists attacks

Not all business interruptions threaten the very existence of the enterprise. However, some industries are far more sensitive to computer system related business interruptions than others.

The "Scale of 9s" Is an Accepted Measure of Computer System Availability

In the field of high-availability engineering, a "Scale of 9s" was first used by respected technologist Jim Gray (formerly of Tandem Computers®) to classify computer system availability into ten levels based on the number of "nines" in the availability equation. As shown in Table 4, a Level 2 computer system is available 99 percent of the time in an average year with an expectancy of one percent unplanned downtime per year. This translates to an expectancy of 87.6 hours of unplanned downtime, which is far greater than a Level 7 system with unplanned downtime of three seconds per year. Typically, fully fault-tolerant computer products, such as the Lightning 9900 V Series and Thunder 9500 V Series systems, exhibit availability of Level 5. Triple redundant computers have availability ratings of 6 or 7. Most competitive products are believed to have availability ratings of Level 4. A Level 4 rating would result in an order of magnitude more unplanned downtime per year for competitive products compared to Hitachi Freedom Storage products.

What is the value of one minute of unplanned downtime to you?

Level	% Availability	% Downtime	Downtime
1	90%	10%	876 hrs.
2	99.0%	1%	87.6 hrs.
3	99.9%	.1%	8.76 hrs.
4	99.99%	.01%	52.56 min.
5	99.999%	.001%	5.25 min.
6	99.9999%	.0001%	31.54 sec.
7	99.99999%	.00001%	3.15 sec.
8	99.999999%	.000001%	.31 sec.
9	99.9999999%	.0000001%	.03 sec.

Table 4 – The "Scale of 9s" is an accepted measure of availability.

The Importance of the Continuous Business Paradigm to Nonstop Operations

Whereas the paradigm for high-availability computing is based on avoiding unplanned interruptions to data availability, Hitachi Data Systems focuses not only on preventing these unplanned interruptions but also on minimizing planned interruptions, such as scheduled downtime for servicing and backup windows. Using a unique, Six-Sigma phased approach, Hitachi Data Systems delivers continuous data availability by applying ISO 9001 type disciplines to threat assessment and risk management. The process begins with a careful analysis of all downtime variables via the Hitachi Data Systems Continuous Business

Planning Questionnaire. In other words, the focus of Hitachi Data Systems is on *continuous business* instead of *business continuity*. This subtle yet important paradigm shift is illustrated in Figure 9.

Hitachi takes a unique approach to high-availability computing in the storage industry. This approach is based on the Six-Sigma standard of quality of operations. First pioneered by Motorola® in the 1980s, Six-Sigma refers to six standard deviations above the mean in sampling of product and process quality. It sets a target of 99.9997 percent defect-free operations for an enterprise, whether for product quality, invoicing accuracy, communications, or any other business process. In short, Six-Sigma means getting everything right. Six-Sigma can be thought of as a process that is similar to those of ISO 9001 in today's modern enterprise. Using Six-Sigma techniques, the Hitachi Data Systems solution requires thorough dedication to all quality processes that affect data availability.

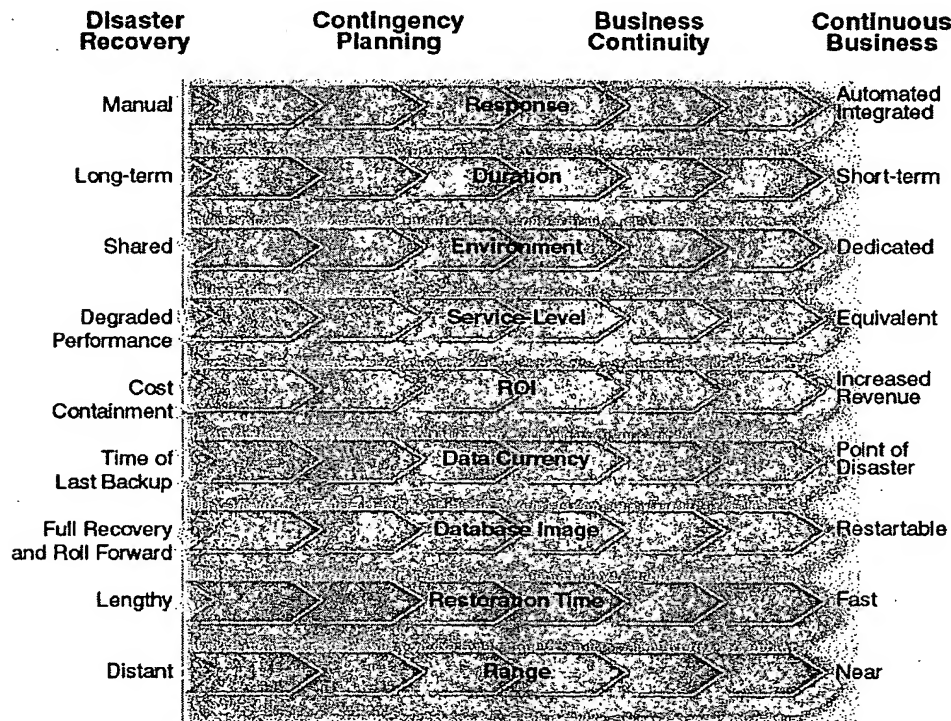


Figure 9 – Hitachi Data Systems changes the paradigm of high-availability computing from business continuity to continuous business.

Implementing the Hitachi Data Systems solution ensures that an enterprise can easily replicate, back up, and manage all information vital to the enterprise business processes.

The Six-Sigma approach reinforces the Hitachi Data Systems reputation for nonstop data availability in the industry, in terms of products, services, and networking. Not only are the company's software solutions for high-availability computing accepted as best of breed by industry analysts and customers alike, but the Hitachi Data Systems tool set for infrastructure assessment and high-availability computing planning is unsurpassed.

The Hitachi Data Systems paradigm of Continuous Business begins by redefining the Scale of 9s. Hitachi looks at the Scale of 9s from an entirely new perspective. By including both planned and unplanned downtime in the availability equation, Hitachi personnel are able to uncover new opportunities for the enterprise to reduce all interruptions – whether they are planned or not. This unique perspective is key to understanding Hitachi Freedom Storage Software Solutions.

There is no doubt that unplanned downtime has an entirely different impact on IT operations than planned downtime. However, a dramatic new picture can be seen when availability is looked at, based on the combination of both planned and unplanned downtime per

The Hitachi Data Systems Continuous Business Paradigm is based on the Six-Sigma standard of quality of operations.

Planned downtime is often an order of magnitude greater than unplanned downtime in terms of real hours when data is not available.

year. Planned downtime can vary from enterprise to enterprise, but on average it is somewhere between 96 hours per year (8 hours per month) and 1,152 hours per year (96 hours per month). Planned downtime is time spent on backups, database reorganizations, bringing new software on line, time for vendor support, etc. Table 5 illustrates two important points. First, whereas the emphasis has traditionally been on unplanned downtime, planned downtime is often an order of magnitude greater in terms of real hours when data is not available. Second, when both planned and unplanned downtime are accounted for, data availability is dramatically lower on the Scale of 9s, as can be seen in the right-hand column of Table 5.

Table 5 – Through the Six-Sigma approach, Hitachi Data Systems attacks both planned and unplanned downtime.

9s Scale	Availability Based Only on Unscheduled Downtime		Availability Based on Scheduled and Unscheduled Downtime			
	Unscheduled Downtime/yr.	Percent Availability	Scheduled Downtime/yr. ²	Total Downtime	Percent Availability	9s Scale
1	876 hrs.	90.0	1152 hrs. (96/mo.)	2028 hrs.	76.84	0.77
2	87.6 hrs.	99.0	1008 hrs. (84/mo.)	1095.6 hrs.	87.49	0.87
3	8.76 hrs.	99.9	864 hrs. (72/mo.)	872.8 hrs.	90.03	0.99
4	52.56 min.	99.99	720 hrs. (60/mo.)	770.8 hrs.	91.20	1.20
5	5.25 min.	99.999	576 hrs. (48/mo.)	576.1 hrs.	93.42	1.34
6	31.54 sec.	99.9999	432 hrs. (36/mo.)	432 hrs.	95.06	1.51
7	3.15 sec.	99.99999	288 hrs. (24/mo.)	288 hrs.	96.71	1.67
8	0.31 sec.	99.999999	144 hrs. (12/mo.)	144 hrs.	98.35	1.83
9	0.03 sec.	99.9999999	96 hrs. (8/mo.)	96 hrs.	98.90	1.89

² In this example, planned downtime is spread over the Scale of 9s chart as shown to reflect the fact that high availability shops tend to have less planned downtime in addition to less unplanned downtime.

Hitachi Data Systems Continuous Business Paradigm, related services, and software solutions provide more key benefits than conventional approaches. These include:

- Conventional hot-site solution alternative
- Dedicated, on-demand, managed rapid-recovery service
- Dedicated testing configuration
- Concurrent change activity support
- Reduced technical complexities
- Reduced communications expense
- Flexible, close secondary facilities at over 500 locations with Hitachi Data Systems partners
- Secure, hardened, and reliable data center facilities
- Comprehensive, highest quality network capability

Hitachi Business Continuity Software Solutions and Services Provide a World-class Selection of Risk Reduction Alternatives for the Enterprise

The Hitachi Data Systems Business Continuity Software Solutions are intended to be implemented to reduce risk of planned and unplanned interruptions. Duplicating and replicating data within a solid worldwide IT architecture, based on best-of-breed products and infrastructure services, are critical techniques for achieving Six-Sigma objectives. These solutions are shown in Table 6. All actions that can move, add, or change existing practices and infrastructure components are considered to reduce the business impact of both expected and unexpected interruptions.

	PRODUCT SUPPORT				OS SUPPORT	
	Thunder 9531V™ Thunder 9532V™ Thunder 9533V™	Thunder 9570V™	Lightning 9900™ Series	Lightning 9900 V Series	Open Systems	IBM S/390®
Hitachi Data Systems Business Continuity Solutions						
Hitachi TrueCopy™		✓	✓	✓	✓	✓
Hitachi ShadowImage™	✓	✓	✓	✓	✓	✓
NanoCopy™			✓	✓		✓
Hitachi Extended Remote Copy		✓	✓	✓		✓
Hi-Track®	✓	✓	✓	✓		✓
Hitachi Dynamic Link Manager™	✓	✓	✓	✓	✓	
Alternative Pathing and Host Fallover	✓	✓	✓	✓	✓	✓

It is important to understand the terminology of Hitachi Freedom Storage Business Continuity Software Solutions as they relate to generic storage terms for copy software. The jargon of copy software alternatives is made even more confusing when traditional backup methods are considered. Advances in technology have allowed new words such as “real time,” “point in time” (PiT) and “snapshot” to creep into the language of enterprise-class storage. Copy products allow an enterprise to replicate, protect, and share data in dynamic new ways. The three main terms used for copy software are:

- Remote Copy
- PiT Copy
- Data Migration

1. Remote Copy

A term that refers to the operations procedure of continuously sending updates to a remote geography in order to provide a time-consistent copy of that data. Synchronous remote copy is typically used over short distances and careful consideration of performance requirements is necessary before implementation of such products as IBM Peer-to-Peer Remote Copy (PPRC), TrueCopy software, and EMC® SRDF™ Synchronous. Asynchronous remote copy uses methods of ensuring data sequencing by timestamping and is used for longer

Duplicating and replicating data within IT architecture are critical techniques for achieving Six-Sigma objectives.

Table 6 – Hitachi Data Systems Business Continuity Solutions.

The jargon of copy software alternatives is made even more confusing when traditional backup methods are considered.

distances. The purpose of remote copy is to protect data in the event of a business interruption to enable business continuity, which can be accomplished with TrueCopy software, HXRC/IBM XRC™, IBM Geographically Dispersed Parallel Sysplex™ (GDPS™) and NanoCopy™ technology.

2. PiT Copy

A process that creates a “static” image of data at a specific time (e.g., backups) is generally referred to as point-in-time (PiT) snapshots, which can be accomplished with Snapshot (StorageTek™), ShadowImage (Hitachi), and TimeFinder™ (EMC).

3. Data Migration

Data Migration software moves data permanently from one storage device to another. This feature is different from data duplication in that at the end of the process there is only one copy of data. The purpose of data migration is to consolidate storage or upgrade to new systems. For example, see Hitachi RapidXchange™ (Chapter 5).

As shown in Table 7, software products provide the functions of data migration, data duplication, and remote copy. These functions are critical to perhaps the two most important business objectives of the enterprise: 1) to maintain business continuity in the face of adversity, such as a disaster; 2) to rapidly deploy new IT applications for business intelligence, data warehousing/OLAP, data center relocation, or new application testing.

Table 7 - Hitachi Freedom Storage software copy products provide remote copy, data duplication, and data migration capabilities.

Hitachi Copy Products	Data Migration	Data Duplication	Remote Copy
Hitachi TrueCopy™ basic – Thunder 9200™, Lightning 9900™ and Lightning 9900 V Series	✓	✓	25 Miles ³
Hitachi TrueCopy asynchronous Extension – Lightning 9900 and Lightning 9900 V Series	✓	✓	Unlimited Distance
Remote Volume Replication – Thunder 9570V™	✓	✓	✓
NanoCopy™ Lightning 9900 and Lightning 9900 V Series		✓	✓
Hitachi Extended Remote Copy (HXRC) asynchronous remote copy – Lightning 9900 and Lightning 9900 V Series		✓	✓
Hitachi ShadowImage™ – Thunder 9200, Lightning 9900 and Lightning 9900 V Series		✓	
Hitachi Data Protection Services – Thunder 9200, Lightning 9900 and Lightning 9900 V Series	✓	✓	✓
Hitachi RapidXchange™ – (HFT) Lightning 9900 and Lightning 9900 V Series		✓	
Tantia® HARBOR® File Transfer – Lightning 9900 and Lightning 9900 V Series Only	✓		

Hitachi TrueCopy Solutions

TrueCopy software can be used to move data locally, between multiple Lightning 9900 V Series systems, between multiple Thunder 9570V Series systems, or between multiple Thunder 9200 storage systems (as shown in Table 8) within a data center, or remotely in geographically separated data centers. There is minimal impact on performance, and no disruption of system operations. Applications continue to run smoothly throughout the remote copy process.

³ While 25 miles is the current practical distance limitation for direct connection using ESCON fibre, this is changing with the advent of new networking devices. Synchronous remote copy may also use telephone networks; however, the performance of the production application must be considered when using synchronous remote copy over long distances.

	PRODUCT SUPPORT				OS SUPPORT	
	Thunder 9531V™ Thunder 9532V™ Thunder 9533V™	Thunder 9570V™	Lightning 9900™ Series	Lightning 9900 V Series	Open Systems	IBM S/390®
Hitachi TrueCopy™						
TrueCopy basic (synchronous)		✓	✓	✓	✓	✓
TrueCopy asynchronous			✓	✓	✓	✓
NanoCopyCopy™			✓	✓		✓

Table 8 – Summary of Hitachi TrueCopy solutions and the products they support.

Hitachi remote copy solutions deliver the choice customers need, regardless of their computing environment. These solutions accommodate all major platforms for both local or remote copy and provide fully consistent data images for disaster recovery for the following supported operating systems:

- Microsoft® Windows NT® 4.0
- Microsoft Windows® 2000
- HP-UX®
- Solaris™ (Sun™)
- AIX® 32 and 64 bit (IBM)
- NetWare® (Novell®)
- Linux® (Red Hat®, SuSE®, Turbolinux®)
- IRIX® (SGI™)
- Tru64™ UNIX® (HP®)
- OpenVMS™ (HP)
- DYNIX/ptx® (IBM)
- NCR® SVR4

Hitachi Data Systems is the only vendor demonstrating proven remote copy support for GDPS, an IBM service offering for system failover, workload balancing, and data mirroring on systems spread across two or more sites, up to 40 kilometers (25 miles) apart. Hitachi Data Systems provides a number of services to help customers plan, integrate, and deploy a GDPS environment based on Hitachi Freedom Storage systems. This is illustrated in Figure 15, later in this chapter, after a discussion of the importance of timestamping.

TrueCopy Basic (Synchronous)

Synchronous remote copy provides the highest degree of data currency and integrity, but can operate only at limited distances. In synchronous communications, the sender and receiver need to synchronize with one another before data is sent. This generally limits the communication between sender and receiver to short distances, typically less than 40km.

Storage system-based synchronous solutions include IBM PPRC, EMC SRDF and TrueCopy basic software as shown in Figures 10 and 11.

Hitachi Data Systems is the only vendor demonstrating proven remote copy support for GDPS.

Synchronous remote copy provides the highest degree of data currency and integrity but can operate only at limited distances.

Figure 10 – TrueCopy basic for open systems environments.





Figure 11 – TrueCopy basic for S/390 environments.



Synchronous remote copy, despite being the most popular form of remote copy technology currently deployed, has a performance deficiency as I/O response times increase with the distance between the primary and secondary sites. This is due to what is known as “propagation delay,” or the fact that, on average, the speed of light takes one nanosecond to travel one foot. This can add up when many miles are involved and synchronous copy techniques are effectively limited to distances of 25 miles or 40 kilometers (10 kilometers for Fibre Channel) depending on performance requirements. These limitations have been solved by Hitachi and are discussed later in this chapter.

TrueCopy (Asynchronous)


TrueCopy asynchronous extension for remote copy is not “real time” like synchronous remote copy, but it can operate at any distance and has the data and transaction integrity necessary for disaster recovery. This is very important to many businesses since many causes of U.S. computer outages affect large geographic areas so that synchronous remote copy at IBM ESCON® distances does not provide real protection. One out of four U.S. computer outages is caused by storms, hurricanes, or floods that span an area greater than that protected by limited-distance synchronous remote copy techniques. Hitachi Data Systems offers a more complete selection of asynchronous copy techniques than any other vendor and provides a number of solutions to address today’s 24/7, anywhere business continuity requirements. Most importantly, other asynchronous remote copies do not have transaction integrity.

In open systems environments, timestamps from Lightning 9900 V Series systems timer are used for data synchronization; in S/390 configurations the host timer or IBM Sysplex Timer® is used. Utilizing reliable timestamps, along with other information created by the primary system(s), allows the Lightning 9900 V Series system to transmit updates directly to the secondary system (without any host intervention) and have the secondary system apply the same consistency group algorithms normally provided by the System Data Mover. This approach allows the enterprise to achieve data integrity in a significantly simpler and less costly hardware-based solution with no impact on server or application performance. Figure 12 shows how Lightning 9900 V Series systems with TrueCopy asynchronous extension utilize timestamps to provide the correct sequencing of “writes” to the secondary volume at the remote site⁴.

Asynchronous remote copy is not “real time” like synchronous remote copy, but it can operate at any distance and has the data and transaction integrity necessary for disaster recovery.

⁴ Lack of individually timestamped I/Os can force an unplanned recovery situation.

Figure 12 – TrueCopy asynchronous extension for open systems.



In the open systems example of TrueCopy asynchronous extension in Figure 13 and the S/390 example in Figure 10, the primary control units are asynchronously connected to control units on the secondary system using reliable inter-system timestamps. Writes (timestamped 1 to 4 in the Figures) are sent asynchronously (i.e., "out of sequence") to the secondary system. These writes are then buffered in the secondary system's cache, queued, sorted by the "timestamp" and written to the correspondent volumes in the same sequence as issued by the primary system over the remote link, checking sequence numbers embedded in records to ensure no records are missing. In this way, I/O consistency is maintained.

There are six major characteristics of TrueCopy asynchronous extension.

1. Data integrity is guaranteed for dependent-write applications.
2. Excellent performance for both long and short distance requirements, due to its asynchronous nature.
3. In S/390 environments, up to four primary systems are supported, and in open systems environments, a one-to-one configuration is supported per consistency group.
4. The maximum distance between the units is 25 miles or 43 kilometers with ESCON connections, or 10 kilometers with Fibre Channel links; for greater distances channel extenders can be used.
5. Control is performed from a graphical user interface (GUI) remote console interface, via a command line interface for open systems (command control software), or via either host PPRC TSO commands or GUI remote console for S/390.

6. Synchronous, asynchronous, open, and mainframe environments may be intermixed within the same Lightning 9900 V Series system.

TrueCopy asynchronous extension in mainframe environments is an excellent product for data relocation and migration activity as well as PiT and real-time copy. TrueCopy asynchronous extension is currently supported in a one-to-one configuration for open systems, and in a four-to-one ratio for mainframes per consistency group as shown in Figure 13⁵. Given that Lightning 9900 V Series systems scale up to 74TB+ raw capacity, it

TrueCopy asynchronous extension in mainframe environments is an excellent product for data relocation.




Figure 13 – TrueCopy asynchronous extension for S/390 environments.

is unlikely that for consistency purposes a customer in open systems would outgrow a one-to-one configuration. In IBM MVS® environments, NanoCopy software is available for larger configurations as is discussed later in this Chapter.

As shown in Figure 14, some competitive products use a technique based on sending changed tracks to the secondary system instead of timestamped I/Os. Users can either specify an "invalid tracks" value, which is the number of changed tracks that can accumulate before that changed data is sent over to the secondary system in full-track images. The problem with this technique is that applying changed tracks (and not individually timestamped I/Os) cannot preserve the original sequence of writes and therefore should not be used for real-time disaster protection or business continuance. Or, in another approach, a table of the changed tracks is accumulated over some period of time (for example, a 24-hour period). At this point, the changed tracks are sent to the secondary location, and when all changed data has been sent, the primary location is quiesced (outage incurred), at which point the two sites are again suspended and the primary site/applications can be restarted. This latter technique is analogous to disk vaulting since it is essentially a point-in-time copy requiring an outage.

⁵ This limitation is only for data that has to be consistent across multiple systems (i.e. Application A and B could run on separate one-to-one configurations if they were separate entities for recovery purposes).

Figure 14 – Not using individually timestamped I/Os can force an unplanned recovery situation.





Figure 15 – TrueCopy software is compatible with IBM GDPS clusters.



NanoCopy

In 2002, Hitachi Data Systems announced the patented NanoCopy feature for the Lightning 9900 V Series systems. With these functions, Hitachi leapfrogged competitors by delivering a solution that removes some of the drawbacks of synchronous remote copy (such as maximum distance and data integrity across multiple systems) without the processor-performance overhead usually associated with asynchronous remote copy implementations like XRC.

TrueCopy asynchronous software for both open systems and S/390 environments is best described as "XRC in a box." The System Data Mover component of XRC contains a number of consistency group algorithms that allow it to maintain data integrity. These algorithms have now been implemented in the microcode of the secondary Lightning 9900 V Series systems.

Working with TrueCopy asynchronous software, NanoCopy technology provides the world's only nondisruptive, I/O- and time-consistent, long-distance remote copy that can be used safely for disaster recovery and PiT copy. NanoCopy software allows "any-to-any" remote copy. This means NanoCopy software provides remote copy with integrity for any number of primary storage systems connected to a single Sysplex Timer, to any number of remote image systems, anywhere in the world, at any time.

Building on TrueCopy asynchronous software, NanoCopy technology uses a cyclical approach to protect any amount of information at unlimited distance, regardless of applications database or data residency (even if resident on multiple controllers). Using a mainframe software routine, a user can set any time for point-in-time copies to be taken, without any outage to the system or impact on performance. Each system suspends the remote data transfers at the same time, creating an I/O-consistent copy (the same "consistency time") across any number of systems. The suspended volume pairs can be resynchronized later.

The major impact of NanoCopy technology on PiT copy technology is that it can create a copy, at any distance, without having to quiesce, shut down, or otherwise interrupt the application that is using (updating) the data. For the first time in the industry, a PiT copy of any amount of data can be created, at any distance, without causing an outage to the application or system and without performance impact. This unique capability vaults Hitachi Data Systems ahead of all other storage vendors for PiT copy capabilities.

NanoCopy technology offers two basic capabilities:

- PiT copy
- Disaster protection through real-time copy

NanoCopy technology comes closest to the perfect solution by allowing customers to create a truly nondisruptive PiT copy. That is, for the first time a storage-based copy solution creates an image of any amount of data without incurring an application or system outage. Furthermore, that copy can be made to any place on the globe. For disaster protection, NanoCopy technology is the only product that solves the extraordinarily complex problem of ensuring that critical data can survive a rolling disaster. With outstanding data integrity, NanoCopy technology not only maintains a copy of critical data at any distance, but also does so using a surprisingly simple and low-cost technique that has minimal impact on critical applications.

NanoCopy technology has the potential for providing many new solutions to the enterprise. It clearly offers a simple and low-cost alternative to existing remote copy technologies without compromising data integrity. It also provides an industry-first capability of taking point-in-time copies without any disruption to applications. Hitachi Data Systems Professional Services consultants are available to tailor, customize, and install the NanoCopy product and its associated software to meet the most rigorous customer requirements. NanoCopy implementation service is discussed further in Chapter 6.

Working with TrueCopy asynchronous, NanoCopy extension provides for the world's only non-disruptive I/O-consistent, and time-consistent long-distance remote copy.

NanoCopy has the potential for providing many new solutions to the enterprise.

Although NanoCopy technology is currently available only on the S/390 platform, Hitachi Data Systems intends to support open systems with the same function, but limited to one primary system and one remote system per consistency group. The one system restriction is required in an open system environment because there is no Sysplex Timer to provide a consistent time source for the entire environment. Instead, open systems use the primary Lightning 9900 V Series system's clock as a time source. The key prerequisite for NanoCopy technology in open systems is a dependable timestamp between primary and secondary systems. This is currently lacking in UNIX and Windows NT or Windows 2000 environments. NanoCopy technology is one of many leading-edge products available from Hitachi Data Systems. NanoCopy technology ensures the safety, integrity, and availability of data in ways that were previously impossible.

The Evaluator Group, a key industry consultant, summarizes the benefits of NanoCopy technology as follows:

With NanoCopy technology, all updates are timestamped and the secondary system uses the timestamps to recreate the original sequence of updates.

"The combination of TrueCopy asynchronous software and NanoCopy technology will provide a relatively simple and cost-effective means of implementing disaster recovery across extended distances. Evaluator Group believes that with these new products, Hitachi has assumed the leadership position in the area of disaster recovery."

NanoCopy technology is built on the current TrueCopy asynchronous technology, which in turn, is built upon the hardware-based concepts of TrueCopy software. All updates are timestamped and sequence-numbered, and the secondary system uses the timestamps and sequence-numbers to recreate the original sequence of updates. NanoCopy technology has multiple components:

- Additional microcode in Lightning 9900 V Series systems to support the two NanoCopy technology commands. These two new commands are NCFREEZE and NCQUERY.
- Host executables (currently in the form of TSO commands on S/390 or RAID Manager for open system environments) to execute the commands to the storage system.
- TrueCopy asynchronous technology and ShadowImage software.
- Automation scripting that allows NanoCopy technology to be used in a disaster recovery environment.

In addition, Hitachi Data Systems Global Solutions Services offers Professional Services that can adapt NanoCopy technology to meet specific customer needs and configurations. The enhancements that NanoCopy technology provides are quite simple. Since all updates are timestamped, NanoCopy technology uses that timestamp to create the I/O-consistent PiT copy. By using NCFREEZE, one of the two new commands available, NanoCopy technology extension is able to "freeze" a second disk image at an I/O-consistent point in time. This process is shown in Figure 16.




Figure 16 – For disaster recovery, NanoCopy triplets ensure there is always a time-consistent set of volumes in the remote system that is not being updated.

The PiT (Point in Time) NanoCopy NCFREEZE command allows customers to specify a future time value at which all primary systems will terminate communication with all secondary systems. For example, if customers wanted to create an I/O-consistent copy of any amount of data at 9:30, they could issue NCFREEZE TIME (9:30) commands to every storage system involved at the primary site at 9:25. By broadcasting this future-time value, they tell each system at which point it is to suspend sending updates to the target. Since all write activity is timestamped, each source system will, upon detecting an update with a timestamp later than 9:30, immediately go into suspend mode. Suspend mode means the primary system would stop sending updates to the secondary system, and begin recording which data is being updated. Put another way, every update up to and including 9:30 has been sent to the target system. All updates occurring after 9:30 are recorded in a bitmap at the source storage system.

This provides the following benefits:

- 1) An I/O-consistent target copy at exactly 9:30 hours
- 2) Protection against a rolling disaster
- 3) The ability to copy only changed data to refresh the target copy asynchronously
- 4) Recovery automation by removing operational decision-making requests
- 5) Reduced recovery time at the remote site
- 6) Logs of all processes
- 7) Reports of all problem areas

Hitachi ShadowImage

ShadowImage software enables open systems users to maintain system-internal copies of Logical Units or Logical Volume Images (LUNs or LVIs) of data for open systems or mainframe environments as shown in Table 9. These copies may be used for purposes such as data backup or data duplication. The RAID-protected duplicate LUNs are created within the same Lightning 9900 V Series, Thunder 9500 V Series, or Thunder 9200 storage systems as the primary LUN at hardware speeds. Once established, ShadowImage software operations continue unattended to provide asynchronous internal data backup. ShadowImage software operations are nondisruptive; the primary LUN of each ShadowImage software pair remains available to all hosts for both read and write operations during normal operations. Usability is further enhanced through a resynchronization capability that reduces data duplication requirements and backup time, thereby increasing user productivity. ShadowImage software also supports reverse resynchronization for maximum flexibility.

Table 9 – Summary of Hitachi ShadowImage solutions and the products they support.

	Product Support				O/S Support	
	Thunder		Lightning		Open Systems	IBM® S/390®
	9200™	9500™ V Series	9900™ Series	9900 V Series		
Hitachi ShadowImage™						
ShadowImage	✓	✓	✓	✓	✓	✓
ShadowImage LVD	✓		✓	✓	✓	✓
ShadowImage FlashCopy Extension	✓	✓	✓	✓	✓	✓

ShadowImage operations can be performed in conjunction with TrueCopy software to provide multiple copies of critical data at both primary and remote sites. ShadowImage software also supports the Virtual LVI/LUN and the ShadowImage LVD (Logical Volume Divider) extension. Hitachi FlashAccess™ features of the ShadowImage FlashCopy extension ensure that all user data can be duplicated by ShadowImage operations.

In a technology exchange with IBM, ShadowImage FlashCopy extension is a software solution whereby the ShadowImage engine responds to all IBM FlashCopy commands from TSO, DFSMSdss and SIBBATCH just like it was IBM FlashCopy®.

ShadowImage software supports the creation of up to nine (three for S/390) RAID-protected image copies from each source volume within a Lightning 9900 V Series, Thunder 9500 V Series, and Thunder 9200 storage systems at hardware speeds, which is shown in Figure 17. There are two key advantages of the Hitachi Data Systems implementation of ShadowImage software:

- RAID protection
- Customer control of copies

The image copy can be of the same or different RAID type. Up to 4,096 image copies are available for Lightning 9980V or Lightning 9970V systems. Once split from the source volume, each copy is independently available for read/write activity.

ShadowImage software further enhances usability through a resynchronization capability that reduces data duplication requirements and thus increases user productivity. When used with TrueCopy software, the duplication process allows users to maintain up to eight image copies in S/390 environments (Lightning 9900 V Series only) and 20 in open systems environments. All copies are RAID-protected.




Figure 17 –
ShadowImage software supports up to nine image copies in open systems and three image copies in S/390 from a single-source volume. The image copies can then be split for other uses such as backup or application testing.

There's nothing more reassuring to IT managers than the knowledge that their company's critical data is fully protected. The high-speed, problem-free data duplication provided by ShadowImage software delivers uninterrupted, 24/7 availability enterprise-wide. Part of Hitachi Freedom Storage software solutions, ShadowImage software makes it possible to duplicate LVIs nondisruptively in the ShadowImage LVD extension. At the same time, primary or secondary remote copy pair volumes can be replicated while maintaining full remote copy capabilities. But it's not just what ShadowImage software does that's important. It's what ShadowImage software enables users to do with it.

ShadowImage Provides the Only Online Backup Solution for Multi-terabyte, Mission-critical Oracle® Databases

Another major benefit of the Hitachi Freedom Storage software solutions is the high-speed, nondisruptive, backup-to-tape solution for users of Oracle databases, as shown in Figure 18. This new capability in data warehousing, which solves one of the most daunting problems facing IT managers today, is the result of an effort to tightly integrate ShadowImage software with the Oracle Database Management System Hot Backup capability. ShadowImage software can provide up to nine copies of a database volume so that the same volume can be used simultaneously for testing, development, data-mining queries, and off-site replication, as well as backups, substantially increasing the value of the solution.

There's nothing more reassuring to IT managers than the knowledge that their company's critical information is fully protected.

The ShadowImage software solution provides a high-speed, nondisruptive, backup-to-tape solution for users of Oracle databases.




Figure 18 –
ShadowImage software can provide high-speed backup of open database systems, such as Oracle.

Hitachi Extended Remote Copy (HXRC)

HXRC is compatible with an IBM Extended Remote Copy (XRC) software, and is based on an IBM asynchronous remote copy technique that is designed to provide complete data integrity for primary and remote systems. This is accomplished through the use of timestamps that allow the asynchronous updates to be formed into consistency groups by the System Data Mover (SDM) host software. Additional support is provided through industry-standard (IBM) compatibility with TSO command support and DASD ERP (error recovery programs/procedures).

HXRC overcomes distance limitations imposed by synchronous remote copy at the expense of potential data loss in the event of a disaster. However, the recovery time to restore full operations is minimized through the use of the SDM and in most cases, the shorter time to accurately recover from a disaster far outweighs the cost of data that might be lost. HXRC is available on both Lightning 9900 V Series and Thunder 9500 V Series systems.

Hitachi Dynamic Link Manager provides load balancing in addition to path failover.

Hitachi Dynamic Link Manager Provides Path Failover and Load Balancing

Dynamic Link Manager is a family of Hitachi-provided middleware software utilities that are server-based, as shown in Figure 19. Dynamic Link Manager software enhances the availability of RAID systems by providing automatic error recovery and path failover from server-to-RAID connection failures. Dynamic Link Manager software provides load balancing in addition to path failover by re-directing I/O activity to the least busy path using complex algorithms.

Just because a system is RAID-protected doesn't mean it is protected against connection bus failures, which is why Dynamic Link Manager software is required for true nonstop operations. This product allows systems administrators to take advantage of multiple paths on Lightning 9900 V Series, Thunder 9500 V Series, and Thunder 9200 storage systems, by adding redundant connections between data servers and RAID systems. Dynamic Link Manager software therefore provides increased reliability and performance. Supported platforms include, IBM AIX, Sun Solaris, Windows NT, and Windows 2000.




Figure 19 – Dynamic Link Manager automatically provides path failover and load balancing for open systems.

Alternate Pathing, Host Failover, and Parallel Database Clusters

Open systems server and third-party software vendors, such as VERITAS® have developed a class of software known as “high-availability middleware” to help reduce downtime by automatically detecting faults and recovering data services on a redundant set of hardware. Without high-availability middleware, time is lost while a fault goes undetected. Once the fault is detected, a diagnose/repair/replace action must take place before data-service recovery can begin. High-availability middleware can begin an automated recovery process immediately on the redundant hardware. The recovery process without high-availability middleware involves time-consuming and error-prone manual operations, which may include resetting the SCSI bus, restarting drivers, reassigning IP addresses, recovering and restarting applications and transactions, and even rebooting.

Alternate Pathing Middleware Switches the I/O Load in the Event of Path Failure

This type of middleware automatically switches the I/O load on a failed primary path to an alternate path on the same host system. Lightning 9900 V Series systems support alternate pathing for AIX 4.2 and above, Tru64, UNIX, HP-UX and through PVlink, Windows NT/2000, Solaris, and DYNIX/ptx. In addition, VERITAS Dynamic Path Management™ (DPM) is supported.

Host Failover

Host failover software supports a cluster of host processors in which one of the hosts automatically takes over the workload of any failed host in the cluster. This “takeover” includes the reassignment of networks and peripherals, as well as the restarting of

The recovery process without high-availability middleware involves time-consuming and error-prone manual operations.

Alternate pathing automatically switches the I/O load on a failed primary path.

One of the hosts automatically takes over the workload of any failed host in the cluster.

applications. Host clustering can also be used to create fault-tolerant workloads and scale-processor capabilities while sharing network and disk resources.

Lightning 9900 V Series and Thunder 9500 V Series systems support all major open-systems clustering schemes, including: HP TruCluster™, HP MC ServiceGuard®, HP MC Lock Manager, IBM RS6000® and SP HACMP®, Windows NT/2000, Microsoft Cluster Server, NCR UNIX SVR4 Lifekeeper, IBM DYNIX/ptx ATAP, and VERITAS FirstWatch® for Solaris.

Clustering middleware supports distributed lock management.

Parallel Database Clustering

This type of middleware is a special version of host-failover middleware, which supports major parallel database servers like Oracle Parallel Server, Informix® XPS, and Sybase® MPP. Clustering middleware supports distributed lock management, a feature that enables parallel database software running on separate cluster nodes to share access to the same database. If one host fails, the other hosts can take over its work. Database clusters allow a customer to grow a database incrementally simply by adding additional nodes. With non-parallel database servers, the server has to be replaced or an additional server with another database instance has to be purchased and installed when the capacity of the original system is exceeded.

Lightning 9900 V Series systems have been certified with MC/Lock Manager for Oracle Parallel Server Version 8, and Sun PDB for Oracle Parallel Server Version 8.

The Hi-Track® Remote Maintenance Tool Ensures the Accuracy of Business Continuity Solutions

The Hi-Track remote maintenance tool for Lightning 9900 V Series systems generates service information messages (SIMs) to identify normal operations (such as TrueCopy software pair status change) as well as service requirements, errors, or failures. SIMs can be generated by the CHIP and ACP microprocessors and by the SVP. All SIMs generated by the Lightning 9900 V Series systems are stored on the SVP for use by Hitachi Data Systems personnel, logged in the SYS1.LOGREC dataset of the S/390 host system, displayed by the Web Console software, and reported via SNMP to the open-system host. The SIM displayed on the Lightning 9900 V Series Web Console PC enables users to remotely view the SIMs reported by the attached Lightning 9900 V Series systems. Each time a SIM is generated, the amber Message LED on the Lightning 9900 V Series control panel turns on. The Hi-Track remote maintenance tool also reports all SIMs to the Hitachi Data Systems Worldwide Support Center.

Backup and Recovery Solutions

4

Overview of Hitachi Freedom Storage™ Backup and Recovery Software Solutions

Backup and restore is the largest cost element in storage ownership and the most important to high-availability computing. In today's global IT environment, it is no longer enough to have a backup copy of one's data; businesses must be able to perform backup and restore in the shortest possible time, and with minimum disruption to information availability. As shown in Table 10, Hitachi Data Systems combines its own leading-edge Business Continuity software solutions with those of industry-leading software partners to address customers' backup/restore requirements with solutions that reduce total cost of ownership (TCO).

Hitachi Data Systems provides backup and restore solutions that reduce the total cost of computing.

	Hitachi Freedom Storage™ Product Support				O/S Support	
	Thunder		Lightning			
	9200™	9500™ V Series	9900™ Series	9900 V Series	Open Systems	IBM® S/390®
Backup and Recovery Solutions						
Hitachi Multiplatform Backup/Restore			✓	✓		✓
Tantia Technologies® HARBOR® Backup/Recovery to HYPERTape			✓	✓		✓
Nondisruptive database backup	✓	✓	✓	✓	✓	✓
Nondisruptive application backup	✓	✓	✓	✓	✓	✓

Table 10 – Summary of Hitachi Freedom Storage backup and recovery solutions and the products they support.

Hitachi Multiplatform Backup/Restore software delivers high-performance, volume-level backup of open systems data. In addition to the Tantia® HARBOR Backup/Recovery and HYPERtape Backup/Restore software packages, Hitachi offers several features and services to achieve this goal. Hitachi ShadowImage™ and Hitachi TrueCopy™ software for IBM® S/390® environments can be used for the creation of a snapshot copy for near-instantaneous backup.

All Hitachi backup and restore solutions are designed to meet the following customer requirements:

- Maximum reliability
- Centralized administration
- File-level backup and restore
- Online backup of popular databases, applications, and messaging services
- Out-of-the-box agents for popular servers, such as Oracle® Server, BMC Software™
- Multiplatform backup to a mainframe or open systems server
- High-performance movement of data
- Ability to "snapshot" the backup source in UNIX®, Microsoft® Windows NT®, Microsoft Windows® 2000, or IBM MVS®
- Ability to simulate or test the recovery scenario

Backup/Recovery Is Critical to High-availability Computing

Backup and recovery is a key component in protecting information. As the number of platforms supported in an enterprise increases, it is increasingly important for an enterprise to choose "best-of-breed" backup and restore solutions. The Gartner Group estimates that Backup and Restore software costs account for 32 to 54 percent of total operation costs. These costs are expected to grow with the increasing need for 24/7 operations, the

What value do you place on the ability to rapidly restore data and recover applications?

implementation of disaster recovery, and the trend towards new applications that generate tremendous amounts of data. Applications such as collaborative computing (Lotus Notes® and Microsoft Exchange), data warehousing/decision support, and business process automation (SAP®, BAAN®, PeopleSoft®) dramatically increase the need for backup and restore. Office email systems are often backed up two or three times a day, and must be recoverable to the document or message level. Business process automation systems are built upon relational databases that provide a single view of data for the entire business from sales and order entry to shop floor control. These requirements place high value on the ability to rapidly restore data and recover applications.

Six Alternative Techniques for Backup and Restore

In order to have consistent, volume-level backup of data, all applications must be stopped.

With the Hitachi Multiplatform Backup/Restore feature enabled, a user can back up any open system to an MVS or VSE system.

IT managers are quickly learning that heterogeneous computing environments present significant challenges.

1. Volume-level Backup/Restore

Of the six major types of backup and restore techniques, Volume-level Backup/Restore is the easiest to do from a hardware and software viewpoint. However, the enterprise must implement controls to ensure the logical and temporal consistency of the data on the volume with the application view of the data. In order to have consistent, volume-level backup of data, all applications must be stopped so that the data in memory and cache can be flushed to the disk. If this volume is part of an application database, it must be coordinated with the backup of the application's logs, indices, repository, updates, exports, and recovery scripts, which are often spread across multiple volumes or generations of volumes. A volume is an MVS concept, and datasets are contained within a volume with a description of the dataset recorded in a volume table of contents (VTOC).

In UNIX or Windows NT/2000, the closest representation to a volume is a SCSI address (Logical Unit or LUN) or mount point (C disk, D disk, and so on).

Unlike MVS, open systems files will often span multiple LUNs and hierarchical directories. Unless the entire file system is backed up at the same time, the restore of volume-level LUN backups is not possible without file-level information from the host file system.

2. Hitachi Multiplatform Backup/Restore Provides Mainframe Benefits

Multiplatform-volume backups between S/390 and UNIX or Windows NT/2000 environments are available with Multiplatform Backup/Restore software on Hitachi Freedom Storage Lightning 9900™ V Series systems. With this feature enabled, a user can back up any open system to an IBM OS/390® or IBM VSE™ system, using any standard utility such as DFDSS or fast dump restore (FDR). Multiplatform Backup/Restore software is the fastest way to back up LUNs to a mainframe and take advantage of the strong infrastructure of mainframe-backup operational procedures.

By using existing backup procedures and Automatic Tape Loading (ATL) environment, Multiplatform Backup/Restore software eliminates the need for special software or procedures. The "any-to-any" connectivity of Lightning 9900 V Series systems enables S/390 users to access these volumes as 3390-3 or 3390-9 devices, making them accessible via standard utilities such as DFDSS or FDR. By working with existing, mainframe-based storage management procedures, Multiplatform Backup/Restore software provides a simple yet effective way for backup/restore of open systems volumes. Multiplatform Backup/Restore software may also be combined with ShadowImage software for near, nondisruptive backup capabilities. This results in the following benefits to open systems backup:

- The use of well-established and disciplined mainframe procedures reduces the chance of human error in backup and restore processing.
- Multiplatform Backup/Restore software capitalizes on high-speed S/390 tape systems.
- Multiplatform Backup/Restore software can even perform a restore of open-system volumes without requiring any host involvement.

IT managers are quickly learning that heterogeneous computing environments present significant challenges. They must rationalize both the equipment and the processes involved in backup and restore operations for critical-business data. This task is relatively easy when data resides on a single host or group of like hosts. But in multiplatform, multiserver environments, data is dispersed on a variety of platforms ranging from S/390 mainframes to all open environments.

Figure 20 demonstrates how open systems volumes can be backed up directly by the S/390 host.




Figure 20 –
Multiplatform
Backup/Restore software
allows UNIX or
Windows NT/2000
volumes to be backed up
to tape using FDR and
mainframe procedures.


File-level backup between platforms requires a software communication utility between a backup agent on a client system and a backup agent on the server.

3. Th Tantia Technology HARBOR File-level Backup and Restore Agent

File-level backup and restore is more difficult to do since it requires the backup software to know and communicate information on the location and extents of a dataset or file. In MVS, this information resides on the disk in catalogs and VTOCs. In UNIX or Windows NT/2000 environments, the information about files and file structures resides in the host file system in mount points and "I-nodes."

File-level backup between platforms requires a communication utility between a backup agent on a client system that can interrogate the file system and a backup agent on the server that records the file-control information for retrieval. As with volume-level backup, additional software and management is required if the file is part of a database or application backup. As shown in Figure 21, HARBOR Backup/Recovery provides a solution for file-level backup between S/390 hosts and UNIX or Windows NT/2000 servers. Usually, the file-level transfer of data between UNIX or Windows NT/2000 clients is done over a network. Hitachi and Tantia have developed a unique agent that provides file-level backup through Lightning 9900 V Series systems and bypasses the network.

Figure 21 – A major benefit of backup/recovery with the HARBOR Agent for Lightning 9900 V Series systems is the ability to restore only single files rather than restoring entire data volumes.



4. Database and Application-level Backup

Databases can be hierarchical (like IMS and VSAM) or relational (like IBM DB2®, Oracle, Informix®, Microsoft SQL Server™, and Sybase®). DB2 is the predominant RDBMS on mainframes, while Oracle is predominant on UNIX and Windows NT/2000 environments. Databases consist of logs, indices, data, and scripts. Relational databases can be further defined by tablespace, which may contain database tables consisting of columns and rows. Databases that are shared require a locking mechanism to maintain logical consistency.

Standalone or "cold" backup of a database requires that the applications against the database are stopped and all the in-flight data is flushed to the disk before backup is started. The applications must remain stopped until the backup is completed and the database is brought back online. In order to reduce this downtime, customers use database utilities to do online or "hot" backups. Online backups still require the applications to be stopped and data to be flushed to the disk. But once this is done to establish a synch point, the applications can resume operations in a "read only" access mode.

Some software utilities provide transaction logging that allows write updates during the online backup. Database vendors are usually the source for these utilities since they require comprehensive knowledge of the database architecture and application. There are some independent software vendors (ISVs), such as BMC, who provide products like SQL-Backtrack™, which can support online backup of multiple vendor database platforms. However, the backup, and, more importantly, the restore of a database requires a database administrator who has a comprehensive understanding of the database and relationships between files. Scripts must be written, tested, and maintained for backing up all the necessary database components.

Online backup of databases requires speed. The longer it takes to do the backup, the more new transactions accumulate in the logs, and the longer it takes to re-synch during the online backup. Hitachi offers several solutions to minimize this time. ShadowImage software can be used in conjunction with products like BMC Enterprise Snapshot™ and SQL products or Oracle Hot Backup Utilities to eliminate any impact or disruption caused by online backup and to take a snapshot copy for backup of mainframe data. The HARBOR file-level backup agent (mentioned earlier in the file-level backup section) can reduce the backup time for databases such as Oracle, Sybase, Informix, and Microsoft SQL Server.

5. Application-level Backup/Restore

Application backup and recovery involves the backup and recovery of logically related groups of database objects across multiple databases or files. Backups must be synchronized to generate a consistent recovery point across multiple files. This requires that all active transactions against target objects be placed on hold. The objects are then switched to read-only mode until the backup is completed. Application backup requires a central management server to control and coordinate backup and recovery. A repository must be maintained on information about all database instances, location of backups, and actions performed against instances. Software like BMC PATROL®, PRAXIS OMNIBACK, and CA™ Unicenter® Framework provide this type of backup.

6. Hitachi Freedom Data Networks™

A major potential benefit of storage area networks (SANs) is LANfree and serverfree backup. Backup windows are the single major issues on the minds of storage executives as determined by a survey by ITCentrix of Fall River, Massachusetts. The Hitachi Data Systems SAN solution is known as Hitachi Freedom Data Networks™. Freedom Data Networks provides an open architecture that leverages SAN technology and offers organizations freedom of choice in deploying data-access and data-sharing capabilities across the enterprise. With Freedom Data Networks, customers gain a powerful new tool that enables the consolidation of servers and storage, increased data availability, centralized storage management, and the ability to back up and migrate data without affecting the performance of enterprise LANs.

Hitachi ShadowImage software can be used in conjunction with products like BMC Enterprise Snapshot and SQL products or Oracle Hot Backup Utilities to eliminate any impact or disruption caused by online backup.

Backups must be synchronized to generate a consistent recovery point across multiple servers.

Backup windows are the single major issues on the minds of storage executives.

The Freedom Data Networks architecture goes beyond SANs.

The Freedom Data Networks architecture exploits advances in servers, storage systems, interconnection devices, network protocols, and network configurations. However, the Freedom Data Networks architecture goes beyond SANs, providing the overall structure for solutions that allow customers to manage their data without being tied to proprietary technology that limits their business options. This gives the enterprise the option to locate storage either inside or outside the data center, wherever it makes the greatest business sense to do so. Moreover, it gives operations personnel the ability to manage a wide variety of server, interconnection, and storage platforms under the Freedom Data Networks methodologies. This provides them flexibility in establishing open system configurations and protecting investment in currently-installed system components.

Within the Freedom Data Networks architecture, SANs provide high-speed Fibre Channel networks for connecting multiple, multi-vendor servers to a pool of multi-vendor storage devices distributed throughout the enterprise. SANs offered by Hitachi Data Systems will support the open-systems standards being developed by the Storage Networking Industry Association.

Industry Alliances are Key to Hitachi Freedom Data Networks Solutions

Key to the delivery of the Hitachi Freedom Data Networks solutions through its direct sales force is a growing number of industry alliances in areas such as networking, communications, and software. Each of these alliances supports Hitachi Data Systems intent to provide fully open solutions based on the use of a wide variety of interoperable components. Alliance partners for solutions developed by Hitachi Data Systems include the following:

- Brocade®, McDATA® and QLogic® for high-performance Fibre Channel switches
- JNI®, JNI Troika™, and Emulex® Corporation for Fibre Channel host bus adapters for Windows NT platforms

Tantia Technologies HARBOR Backup/Recovery and HYPERTape

Hitachi Data Systems addresses the needs of both MVS and LAN customers with backup and restore solutions that are scalable to the enterprise level. Intended respectively for MVS and LAN environments, HARBOR Backup/Recovery (HBR) and HYPERTape support a full menu of open systems clients, mainframes, and agents for both online and offline backup of databases.

VERITAS® NetBackup™ Offers Backup and Recovery Solutions for Heterogeneous Environments

Scalable from the desktop to the data center, NetBackup delivers quick and reliable backup and recovery that spans terabytes to petabytes in size. The VERITAS NetBackup Array Integration option and VERITAS NetBackup ServerFree Agent are tightly integrated with ShadowImage software and the Hitachi e-Copy feature, allowing administrators to automate split-mirror backups on Hitachi Freedom Storage systems with zero impact on applications availability. The result? A significant reduction in the use of server-based resources traditionally required for backups.

* Hitachi Data Systems is a reseller of HARBOR and HYPERTape backup/restore software packages in the U.S. and Canada.

Data Movement Solutions

5

Hitachi Data Systems developed its Data Movement Solutions to provide heterogeneous IT environments with the critical ability to transfer copy and share large amounts of data between heterogeneous platforms and applications. This capability is key to business success in today's global information processing environments, where corporate data must be readily available to all worldwide users, regardless of their location or the nature of their computing environment. Data from heterogeneous systems needs to be exchanged regularly, either for data synchronization between discrete applications or, increasingly, for populating data warehousing systems.

Table 11 lists the Data Movement Solutions Suite that supports the Hitachi Freedom Storage™ Lightning 9900™ V Series and Thunder 9200™ storage systems.

	Hitachi Freedom Storage™ Product Support				O/S Support	
	Thunder		Lightning			
	9200™	9500™ V Series	9900™ Series	9900 V Series	Open Systems	IBM® S/390®
Data Transfer Solutions						
Hitachi RapidXchange™			✓	✓		✓
Hitachi Data Migration Services (Chapter 6)			✓	✓		✓
Tantia Technologies® HARBOR® File Transfer				✓		
Hitachi e-Copy				✓	✓	

Table 11 – Data Movement Solutions provided by Hitachi Data Systems.

Hitachi RapidXchange™ Software Enhances Decision Support Productivity while Improving Operations Efficiency

The ability to share and move data between heterogeneous environments, at very high speeds, while minimizing the impact on the overall throughput of network communication links is a significant asset. It is especially useful for data warehousing and data mining, where vast amounts of data must be moved quickly from mainframe OLTP systems to UNIX® or Microsoft® Windows® environments. In some industries, the time it can take to refresh a data mart or data warehouse can define a competitive edge.


Hitachi RapidXchange software does not need additional software on an IBM MVS® host. RapidXchange software provides maximum platform flexibility and transfer speed for the quick and reliable data transfer. Any data that can be unloaded, created, or copied as a sequential dataset in MVS can be read directly by a UNIX server using RapidXchange software. Multiple, disparate hosts (for example, HP®, Sun™, and Windows) can access a single copy of data. RapidXchange software is especially useful for data warehousing and data mining tasks, where vast amounts of data need to be moved from mainframe OLTP systems to a UNIX environment (such as data warehouses or data marts) in a very short time. RapidXchange software delivers this capability without requiring networking resources or involving intermediate media, such as tape. High-speed data transfer is achieved without putting additional data load on the network infrastructure or tape transport equipment. Also, RapidXchange software has ultrahigh-speed transfer rates (when used in conjunction with Hitachi FlashAccess™ software), which delivers the fast response demanded by today's widespread data-warehousing and data-mining applications.

RapidXchange software provides for much higher Fibre Channel effective data transfer rates than today's networks can accommodate. Lightning 9900 V Series system users can take advantage of even faster data transfer rates by using RapidXchange software in conjunction with FlashAccess software, which maintains the intermediate volume in the Lightning 9900 V Series system's cache. Using cache as a buffer eliminates any physical disk

Hitachi RapidXchange software offers high-performance and high-reliability data sharing between heterogeneous host platforms such as IBM® S/390®, UNIX, or Windows.

Figure 22 –
RapidXchange software
quickly converts
mainframe EBCDIC
files to open systems
ASCII files for
nondisruptive
population of data marts
and data warehouses.

RapidXchange software
transfers data at the
speed of Fibre Channel.



access during the write/read process and enables data exchange at channel speeds. Data can be transferred from a UNIX or Windows system to an MVS system by using the same procedure in the reverse direction. This process is illustrated in Figure 22.

RapidXchange software eliminates many of the constraints associated with sharing or moving data between heterogeneous systems. Sequential files on common storage volumes allow large amounts of data to be shared or moved. Using RapidXchange software, you can "share" sequential datasets between MVS and open-systems platforms.

RapidXchange software provides two options for accessing data from open systems platforms:

- File Access Library
- File Conversion Utility

The UNIX flat file to
receive the transformed
data can be resident on
any UNIX volume,
including non-Hitachi
storage.

File Access Library API

The RapidXchange File Access Library provides basic API functionality in the form of open, read/write, and close commands. UNIX or Windows applications can use this API to read MVS data directly from a RapidXchange disk volume and then process that data. In many cases, existing programs that currently use tape input can be modified easily to read that data directly from the MVS system, thereby eliminating the MVS-to-tape-to-UNIX process.


File Conversion Utility

A second option is the Hitachi graphical user interface (GUI)-based File Conversion Utility (FCU). Supplied as a standard component of RapidXchange software, the FCU furnishes a GUI-based display of the intermediate MVS volume. From this display, individual sequential datasets can be selected, converted to ASCII, and transferred as a simple UNIX or Windows flat file. Transformation options include converting data from EBCDIC to ASCII or padding out variable-length records. The UNIX or Windows flat file to receive the transformed data can be resident on any UNIX volume, including non-Hitachi storage.

Tantia Technologies® HARBOR® File Transfer Increases Speed and Automates Performance of Large File Transfers

Running with RapidXchange software, HARBOR File Transfer adds automation to the process of transferring large data files, at high channel speeds, in either direction between open systems and mainframe servers. It can be thought of as an ultra-fast FTP transfer, except data is transferred over an IBM ESCON® link instead of Ethernet. HARBOR File Transfer enables UNIX and Windows applications, such as enterprise resource planning and decision support systems to work with up-to-the-minute data extracted from IBM S/390® environments. Multithreading capabilities and Fibre Channel speeds of 100MB/sec support HARBOR File Transfer's high-speed data movement. After automatically breaking large data files into more manageable pieces, HARBOR File Transfer offers increased transfer speed by directing data in multiple streams through the Lightning 9900 V Series system(s).

A HARBOR File Transfer agent running on the open systems client platform sets up control information (including authorization and file lists) and sends it to the MVS host



HARBOR File Transfer adds automation to the process of transferring large data files at high channel speeds.

Figure 23 – HARBOR File Transfer runs on RapidXchange software and transfers large data files at channel speeds in either direction between open systems and mainframe hosts.

HARBOR File Transfer is more efficient than transfer over the much slower network lines used in conventional solutions.

over the network as shown in Figure 23. When MVS returns acknowledgment, the client sends the corresponding file data to the host through RapidXchange software. This reduces the traffic on the network, and minimizes the backup time by transferring data at ESCON, IBM FICON™, and Fibre Channel speeds rather than at network speeds.

This is more efficient than transfer over the much slower network lines used in conventional solutions. As shown in Figure 23, HARBOR File Transfer employs the network to transfer only “control information” between open systems and mainframe environments, so they remain more fully available for regular network activity. An important timesaving feature of HARBOR File Transfer is its capability, in the event of a transmission failure, to restart from the point of failure. This means that transfers can be continued from the point in time that a failure occurred, rather than from the very beginning of the transfer process. This saves a tremendous amount of time compared to FTP network alternatives. With the added features of alternate path routing (through the network or through Lightning 9900 V Series systems) and retry/reconnection logic, HARBOR File Transfer brings heightened reliability and fault tolerance to the file transfer process.

For state-of-the-art management, HARBOR File Transfer offers an easy-to-use GUI that monitors transfer rates, active transfers, logging, and system controls, and simplifies file selection and batch file creation. HARBOR File Transfer's GUI allows both system administrators and end users to share mainframe and client system data and perform backups and restores. Efficient file transfer has become a major challenge for many organizations. Tape technology often requires manual intervention (and the risk of human error) to move data between separate systems. Today's network congestion often does not provide the performance necessary for the transfer of large volumes of data in a short time. This is due to the conflict for transmission with other network traffic and network protocol overhead.

Hitachi e-Copy Provides LANfree and Serverfree Data Movement

A common function for Hitachi Data Movement Solutions is for LANfree and Serverfree disk-to-disk or disk-to-tape data transfer for backup purposes. The Hitachi e-Copy feature, when used with applications on the open-systems host server, such as the VERITAS® NetBackup™ series of products, and with Hitachi ShadowImage™ software, provides serverfree backup solutions between the Lightning 9900 V Series system and backup devices, such as tape and disk devices. The e-Copy feature enables nondisruptive backup directly from disk to tape (or disk) in storage area network (SAN) environments, eliminating server CPU and I/O overhead during movement of data and decreasing the time required for backup. Users can perform copy (backup) operations on the data (in units of block) stored on the multiplatform Lightning 9900 V Series system via Fibre Channel interfaces directly to the backup devices.

The e-Copy feature supports the SCSI Extended Copy command issued from the host server to the Lightning 9900 V Series system. The Lightning 9900 V Series system receives the e-Copy commands issued by the server, and then copies the data directly to the specified backup device. E-Copy operations are configured using the LUN Manager Remote Console – Storage Navigator software. To implement e-Copy operations, a backup application on the host server must issue the e-Copy commands to the Lightning 9900 V Series system. The LANfree and serverfree data movement capability of the e-Copy is shown in Figure 24.




Figure 24 – Working with the VERITAS NetBackup series of products, the Hitachi e-Copy feature provides LANfree and serverfree data movement.

Professional Services and Support

6

Hitachi Data Systems Is Consistently Ranked Number One in the Industry

In numerous independent surveys on IT services organizations, Hitachi Data Systems continually wins the highest ratings in terms of overall customer satisfaction. "Service Responsiveness" is the key Hitachi Data Systems characteristic that allows the world-renowned Hitachi Data Systems service and support organization—Global Solution Services—to ensure that Hitachi products operate at peak performance to complement all hardware and software in the enterprise.

Professional Services Overview

As new technologies gain acceptance, companies must decide on long-term plans and implementation schedules that cause the least disruption to business. It takes time to implement any large-scale technological change. The transition to new network topologies will see the coexistence of distributed and legacy systems, as well as SCSI and Fibre Channel on storage area networks (SANs), storage wide area networks (SWANs), IBM® ESCON®, and IBM FICON™. Hitachi Freedom Storage™ provides the comprehensive connectivity, management, and availability capabilities needed to handle this transition. Hitachi built-in product strengths are bolstered by the Hitachi Data Systems Global Solution Services organization to ensure the optimal operation of hardware, software, and middleware for the enterprise.

Global Solution Services employs Professional Services teams specializing in infrastructure, hardware, software, and storage management services that provide a vendor-independent view of IT architecture. This allows a sharp focus on ways to streamline operations, costs, and interoperability. Hitachi Data Systems also excels at helping customers chart both the strategies and timelines necessary to remain productive and competitive. Whether an enterprise needs assistance with SANs, business continuity consulting and implementation, DFSMS performance/capacity issues, migration planning, decisions about platforms and architectures, or maximization of IT investments, Hitachi Data Systems has the expertise and the resources to guide an enterprise toward the best business solution. A few of the many Global Solution Services offerings are highlighted here.

*Hitachi Data Systems
Global Solution Services
delivers comprehensive
connectivity,
management, and
availability capabilities
through Professional
Services teams.*

Data Protection Services

Our Professional Services team is trained in architectural analysis, configuration planning, and enterprise assessment to deliver Data Protection Services. The team uses structured methodologies that promote consistent results. By applying world-class software solutions, the Professional Services team implements essential copy solutions by moving data in order to protect it, and employs software in all three categories of data movement:

- 1) Data migration
- 2) Real-time copy
- 3) Point-in-time copy

The team will perform data relocation and migration tasks that establish rock-solid backup and disaster recovery copy facilities to keep a business running in the event of a man-made or natural disaster. Often, the need for an enterprise to allocate staff for these tasks is eliminated.

With Data Protection Services, an enterprise can accelerate its return on existing investments, streamline data management processes, eliminate redundant software and costs, and reduce time to recovery.

As a first step in a typical engagement, an Engagement Manager works with the enterprise to determine requirements based on the use of a proven Data Movement questionnaire and the presentation of examples of prior engagements. Disaster recovery, backup, data warehousing, business continuance, testing and development, and data center consolidation are the types of projects that most frequently use this service. Then the team reviews system, network, storage components, and phases of solution delivery with the client. After a clear understanding of requirements is reached, a detailed professional services proposal is presented and reviewed with the enterprise. Hitachi Data Systems uses “best of breed” consulting techniques and has an excellent reference list of satisfied Global Solution Services clients who may be contacted prior to beginning of the engagement. Examples of the Global Solution Services suite of data protection services are discussed below:

Data Protection Services Lab

Hitachi Data Systems has designed a laboratory proving ground for developing the tools and expertise an enterprise needs to fully enable the functions of Hitachi Freedom Storage systems and software. The Data Protection Services Lab is dedicated to resolving software and hardware issues that affect storage system consolidation, data center productivity, and business continuity.

Remote Copy Assessment and Implementation Service

This service assists the enterprise in implementing a remote copy process for either disaster recovery or for the rapid deployment of new IT systems, such as for relocation of a data center, population of a geographically dispersed data mart, or testing new applications. Global Solution Services provides expert consultants to assist in assessment of processes and procedures to ensure an optimal implementation of the remote copy process. Global Solution Services also offers services for developing remote testing at a disaster recovery hot site to validate the proper configuration of the remote copy process in terms of hardware, software, and processing method. Enterprise IT professionals are trained to address remote copy as a strategic tool of the data center to meet the objectives of business continuity, disaster recovery, or rapid deployment of new IT systems.

NanoCopy™ Implementation Service

NanoCopy Implementation Service lets an enterprise make point-in-time copies without disturbing critical applications or causing any disruption to end-user operations. It is the only completely nondisruptive, long-distance remote copy technique that can be used for disaster recovery with confidence. Global Solution Services experts are available to help customize and install this advanced alternative to existing remote copy technologies. This service helps the enterprise adapt the NanoCopy technology to the most rigorous requirements, ensure maximum operations efficiency and availability, and shorten enterprise recovery time dramatically in the event of a disaster.

Continuous Availability Service

Keeping critical systems available at a level 4 or 5 on the Scale of 9s is no easy task. It involves the best of hardware, software, and operations practice. Global Solution Services experts review the entire backup and recovery software portfolio, operations procedures, and hardware configurations of the enterprise. Weaknesses that threaten the IT environment are then systematically analyzed and eliminated. With this service, Global Solution Services consultants employ the best consulting practices in both open systems and IBM® MVS™ application backup and recovery (both onsite and offsite) to conduct the assessment. As a first deliverable, a baseline will be created and presented for analysis of potential strengths and weaknesses of the current processes, equipment, software, and practices. Recommendations are then discussed to improve availability and reduce the risk of data loss.

Global Solution Services provides expert consultants to assist in assessment of processes and procedures.

Global Solution Services consultants employ best consulting practices in both open systems and MVS application backup and recovery.

Glossary of Terms

Alert

A message or log that a computing element generates as the result of an error event collection and analysis. An alert indicates that there is a need to perform some service action and can be sent by a variety of methods to operations personnel.

API

Application Interface or API is a set of calls that allow software developers to interface to a specific program.

Asynchronous

Asynchronous communication (as in asynchronous remote copy) occurs when the transmission of data between two devices is not synchronized with a clocking scheme or other technique. The sender can send data at any time and the receiver can accept information when it becomes available. Synchronous communication is an exactly timed stream of bits when the start of a character is located by using a clocking mechanism such as bipolar encoding. Asynchronous and synchronous transmissions are used extensively in the mainframe terminal environment.

Availability

In computer science, availability refers to the degree to which a system or resource is capable of performing its normal function. Availability is measured in terms of Mean Time Between Failure (MTBF) divided by MTBF plus the Mean Time to Repair (MTTR). The availability equation is expressed as follows:

$$\text{AVAILABILITY} = \text{MTBF} / (\text{MTBF} + \text{MTTR}).$$

For example, a server fails on average once every 5,000 hours and takes an average of two hours to diagnose, replace faulty components, and reboot, would have an availability rating of $5,000 / (5,000 + 2) = 99.96\%$. This would correspond to a Level 3 rating using the Scale of 9s as discussed in Chapters 1 and 2.

Business Continuity Planning (BCP)

An "umbrella" term covering both disaster recovery planning and business resumption planning. See also Disaster Recovery.

Business Impact Analysis (BIA)

The process of analyzing all business functions and the effect that a specific disaster may have upon them.

Business Interruption

Any event, whether anticipated (i.e., public service strike) or unanticipated (i.e., blackout) that disrupts the normal course of business operations at a corporate location.

Cache

Cache (pronounced cash) can be either on-chip memory circuits in a microprocessor (e.g. L2 processor cache), a reserved section of main memory (system or server cache), or an independent high-speed disk storage device (e.g. a Web cache). Two types of caching are commonly used in personal computers: memory caching and disk caching. Disk caching can dramatically improve the performance of applications, because accessing a byte of data in RAM can be thousands of times faster than accessing a byte on a hard disk. When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate.

CIM

The Common Information Model, or CIM, is an open data model that provides for the collection, storage, and analysis of management data for systems, networks, applications, and storage in a common format. CIM is being developed by the Distributed Management Task Force (DMTF) and provides a means for vendors to present management data about their products, independent of their underlying interfaces and data formats.

Client/Server Architecture

Client/Server Architecture is a network architecture in which each computer or process on the network is either a client or a server. Servers are powerful computers or processes dedicated to managing disk drives (file servers), printers (print servers), or network traffic (network servers). Clients are PCs or workstations on which users run applications. Clients rely on servers for resources, such as files, devices, and even processing power.

Cold site

An alternate facility that is void of any resources or equipment except air-conditioning and raised flooring. Equipment and resources must be installed in such a facility to duplicate the critical business functions of an organization. Cold sites have many variations depending on their communication facilities, UPS systems, or mobility. See also Shell site, Recovery site, Alternative site.

Common Information Model (CIM)

The Common Information Model (CIM) is an object oriented emerging standard to the central management of systems and networks across multiple organizations with heterogeneous systems and software. CIM is sponsored by the Distributed Management Task Force which receives contributions from participating technology companies. It is also supported by the SNIA.

CVS

Customized Volume Size for open systems. A feature in the Lightning 9900™ V Series storage systems which enables a user to define LUNs that are non-standard size.

DASD

An IBM® coined acronym that means Direct Access Storage Device, i.e. disk drives.

Data Availability

Data availability refers to the degree to which a computer system is capable of providing data to users. See also Chapter 3 and Availability.

Data Copy

A term that refers to remote copy, data duplication, and data migration. See also Chapter 3.

Data Duplication

Software that duplicates data as in remote copy or point-in-time (PiT) snapshots. Data duplication is differentiated from data migration in that with data duplication at the end of the process there are two copies of data and with data migration there is only one.

Data Migration

Software that migrates data from one storage device to another. This feature is different from data duplication in that at the end of the process there is only one copy of data.

Disaster

Any event that creates an inability for an organization to provide critical business functions for an undetermined period of time.

Disaster Recovery

Disaster Recovery is the profession that plans to avoid disasters and to restore services after a disaster. It also indicates the level of preparedness to respond to an interruption in services by implementing a disaster recovery plan to restore an organization's critical business functions is discussed in Chapter 3.

DR

Disaster Recovery.

Downtime

A planned or unplanned interruption in system availability. Planned downtime is usually for scheduled system maintenance and unplanned downtimes usually include business interruptions or disasters. See also Business Interruptions, Disaster Recovery, Business Continuity Planning (BCP), and Availability.

ECC

Error correction code.

Electronic Vaulting

The transfer of data to an offsite storage facility via a communication link rather than via portable media. Typically, electronic vaulting is used for batch or journaled updates to critical files to supplement full backups taken periodically.

ESA/390®

Enterprise Systems Architecture/390®. Often abbreviated S/390®, ESA/390 is an IBM architecture for mainframe computers and peripherals. Processor systems that follow this architecture include the IBM ES/9000® family. In 2000, IBM changed the server series name to eServer zSeries™.

ESCON®

Enterprise Systems Connection architecture is an IBM mainframe ESA/390® computer peripheral interface or connection between two mainframes for data exchange. The I/O interface utilizes ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric. ESCON is based on networking technology. ESCON provides direct channel-to-channel connections between mainframe systems over fiber-optic links at distances up to 43 kilometers or 25 miles. ESCON also provides a way for communication controllers and other devices to share a single channel to a mainframe.

ESCON Director

An I/O interface switch that allows the interconnection of multiple ESCON interfaces in a distributed-star topology.

Ethernet

A Local Area Network (LAN) protocol developed by Xerox® in cooperation with Digital Equipment and Intel in 1976. Ethernet supports a star or bus topology and supports a data transfer rate of 10 megabits per second or 10 Mbit/sec. The Ethernet specification formed the basis of the IEEE 802.3 standard, which specifies the physical and lower software layers. Ethernet uses the CSMA/CD access method for handling simultaneous demands and is one of the most widely implemented LAN standards. Ethernet is also known as 10BaseT. See also Fast Ethernet, Gigabyte Ethernet.

Fabric

A fabric is one of three Fibre Channel topologies. In a Fabric topology Node Ports (N_Ports) are connected to Fabric Ports (F_Ports) on a switch. See also Switch.

Failover

Host, host bus adaptor, cable, or controller failover is the routing of all transactions to a second host or controller when the first component fails.

Fast Ethernet

Fast Ethernet or 100BaseT, defined by the IEEE 802.3 committee, provides a 100Mbit/sec standard that is compatible with existing 10BaseT installations, preserving the CSMA/CD media access control (MAC) protocol.

Fast write

A write operation at cache speed that does not require immediate transfer of data to a disk drive module. The system writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for de-staging (writing to disk). Fast write reduces the time an application must wait for the I/O operation to complete.

FC

An acronym for Fibre Channel.

FC-AL

Fibre Channel Arbitrated Loop is the most dominant of the three topologies of Fibre Channel. Loops are a cost-effective way of connecting up to 127 ports in a network without the need for a switch. See also Fabric and Switch.

FCP

An acronym for Fibre Channel Protocol – an ANSI standard covering Fibre Channel protocol for SCSI. See also Fibre Channel Protocol.

Fibre Channel

Fibre Channel is an ANSI standard designed to provide high-speed data transfers among workstations, servers, desktop computers, and peripherals. Fibre Channel makes use of a circuit/packet switched topology capable of providing multiple simultaneous point-to-point connections between devices. Fibre Channel is widely deployed in SAN implementations today. Standards for Fibre Channel SANs are being worked on by the Storage Networking Industry Association (SNIA). The technology has gained interest as a channel for the attachment of storage devices, but has limited popularity as high-speed networks interconnect. Fibre Channel can be deployed in point-to-point, arbitrated loop (FC-AL), or switched topologies. Fibre Channel nodes log in with each other and the switch to exchange operating information on attributes and characteristics. This information includes port names and port IDs and is used to establish interoperability parameters.

Fibre Channel Protocol (FCP)

Fibre Channel Protocol is an ANSI standard covering Fibre Channel protocol for SCSI.

FICON™

The Fibre Connector channel (FICON) is IBM's trademarked channel for zSeries® or S/390® processes in SAN environments.

File Backup

The practice of copying a file that is stored on disk or tape to another disk or tape is referred to as file backup. This is done for protection case the active file gets damaged. Backup is considered "local copy", as opposed to "remote copy." See also Remote Copy.

Gigabit Ethernet

Provides a standard that supports data transfer at 1000 Mbit/sec. Gigabit Ethernet is also called 1000BaseT Category 5 (copper wire) or 1000BaseX (fiber optic). There is a 10,000BaseT version of the Ethernet standard that will be widely available by 2003.

GUI

GUI is an acronym that refers to a Graphical User Interface, which is the software that controls the screen presented to a user in a computer application.

HARBOR® File Level Backup/Recovery (HBR)

HARBOR File Level Backup/Recovery is a software utility from Tantia Technologies® that provides for network or channel-based backup/restore of open systems client files to an OS/390® host using a common graphical user interface. Online, non-disruptive backup/recovery of popular database, mail, and Enterprise Resource Planning (ERP) applications are optional. Client support is available for Microsoft® Windows®, OS/2®, UNIX®, Novell® NetWare®, and IBM® VMS® host platforms. High-speed data transfer is also available using Hitachi RapidXchange™ and other channel-based technologies.

HARBOR File Transfer

This software utility provides automatic, reliable, and secure data transfer between OS/390 and open systems hosts using high speed IBM ESCON®/FICON™ channels or network communications. HARBOR File Transfer can be thought of a very high speed FTP that does not use a network. High performance is enabled by allowing multiple file transfers to run concurrently. Reliability is enhanced with retry logic, alternate path routing, and restart from point of failure. Ease of use is provided by a Sun™ Java™-client interface, a system monitor, scheduling support, and remote control of client functions from IBM OS/390®.

Hi-Star™ architecture

At the heart of the Lightning 9900™ V Series revolutionary design is the Hi-Star architecture, which provides multiple, redundant, non-blocking paths between the storage ports, multiple cache nodes, and multiple disk Array Control Processors (ACPs). Each path runs at 100MB/sec, permitting Lightning 9900 V Series systems to scale up to a total bandwidth of 15.9GB/sec. This is six-to-eight times the internal bandwidth of shared-bus based storage systems. For redundancy and performance, Hi-Star consists of four cache switches that are cross-connected to four cache modules and two control memory banks. These control memory banks contain addressing and control information and are also cross-connected to front-end storage ports and back-end disk array ports.

Hitachi CruiseControl™

CruiseControl automatically monitors, analyzes, and *moves logical volumes* to eliminate "hot spots" within a Hitachi Freedom Storage™ Lightning 9900™ V Series storage system and provides load balancing to maintain predetermined performance levels.

Hitachi Data Systems Migration Service

The Hitachi Data Systems Migration Service (HDmS) is a Hitachi Data Systems professional service that helps users migrate data from existing systems to newly installed systems while minimizing the impact on mission-critical applications. HDmS features a four-phase approach that includes assessment, planning, migration, and post-migration support.

Hitachi Dynamic Link Manager™

Hitachi Dynamic Link Manager is a family of software utilities that is server based and enhances RAID systems by providing automatic failover and load balancing from server-to-RAID channel connection failures. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900 V Series system by adding redundant SCSI connections between data servers and RAID systems. Hitachi Dynamic Link Manager therefore provides increased reliability and performance. Supported platforms include AIX®, Sun Solaris™, and Microsoft® Windows NT®/Windows® 2000.

Hitachi e-Copy

E-Copy is a command in the SCSI 3 specification that allows for the transfer of data without the involvement of the server. The Lightning 9900™ V Series systems implement this command, which is used by programs such as VERITAS® NetBackup™ to enable serverless and LANfree backup.

Hitachi FlashAccess™

FlashAccess software allows specified (usually high access) data sets to be "pegged" or permanently placed in cache memory so they are not managed by the data movement algorithms of the Hitachi Freedom Storage™ system. The FlashAccess feature in the Lightning 9900™ V Series systems can be used for either S/390 or open systems. FlashAccess is a software utility in the Hitachi Resource Manager™ suite that allows the creation, deletion, and monitoring of data managed by the FlashAccess software. See also Hitachi Resource Manager.

Hitachi Freedom Data Networks™

Hitachi Freedom Data Networks provides an open architecture that leverages SAN, NAS, and IP technology and offers organizations the freedom of choice in deploying data-access and data-sharing capabilities across the enterprise. With Freedom Data Networks, customers gain a powerful new tool that enables the consolidation of servers and storage, increased data availability, centralized storage management, and the ability to back up and migrate data without affecting the performance of enterprise networks.

Hitachi Graph-Track™

Graph-Track is a software utility in the Resource Manager suite that enables a robust set of system and network management utilities and provides graphical reports for Lightning 9900 V Series performance, availability, and configuration management.

Hitachi LUN Manager

Hitachi LUN Manager is a software utility in the Hitachi Resource Manager™ suite that allows for complete systems management of LUNs. See also LUN, and Hitachi Resource Manager.

Hitachi Multiplatform Backup/Restore

Hitachi Multiplatform Backup/Restore is a channel-based backup/restore of open systems volumes using standard mainframe utilities, which leverage current investment in mainframe hardware, software, skills, and procedures. Hitachi Multiplatform Backup/Restore provides high-performance, high-bandwidth capabilities. One copy of software resides on each controller.

Hitachi Multiplatform Resource Sharing

Hitachi Multiplatform Resource Sharing is built right into the Hitachi Freedom Storage™ Lightning 9900™ V Series systems. With Hitachi Multiplatform Resource Sharing, Lightning 9900 V Series system users can share resources between UNIX®- or Microsoft® Windows NT®/Windows® 2000-based servers and IBM® S/390® mainframe platforms. Sharing resources across heterogeneous platforms lowers total cost of ownership, provides a centralized point for data management, and simplifies the management of heterogeneous systems.

Hitachi Priority Access

Hitachi Priority Access allocates bandwidth on the basis of quality of service (QoS) requirements either at the physical port level or at the Host Storage Domain level within a physical port. See also Host Storage Domains.

Hitachi RapidXchange™

Hitachi RapidXchange provides for the file conversion and exchange of data between S/390 and open systems hosts. Hitachi RapidXchange provides file access APIs for open system hosts and runs on the following host servers: AIX, Sun Solaris, HP/UX®, Compaq® Tru64™ UNIX, Sequent® DYNIX/ptx®, SGI™ IRIX®, NCR® UNIX SVR4®, and Windows NT/2000.

Hitachi Resource Manager™

The Hitachi Resource Manager is a comprehensive suite of management software that brings together Hitachi Graph-Track™, Hitachi Virtual Logical Volume Image (VLVI) Manager, Hitachi FlashAccess™, and Hitachi LUN Manager software into one complete package.

Hitachi SANTinel™

Hitachi SANTinel software controls host access to Hitachi Freedom Storage 7000E or Lightning 9900 V Series LUNs in SAN environments.

Hitachi ShadowImage™

Hitachi ShadowImage is a firmware-based software copy utility that uses command-line interfaces to create up to ten copies of a volume within one Hitachi Freedom Storage™ Lightning 9900™ V system, or up to 20 across multiple Lightning 9900 V Series systems. Graphic or command-line interfaces control data replication and fast resynchronization of logical volumes. ShadowImage is available for open systems or IBM® S/390® environments. ShadowImage also works in concert with Hitachi TrueCopy asynchronous for S/390 to provide additional copies in another system. ShadowImage for open systems is a point-in time (PIT) copy capability that allows the data replication and fast resynchronization of logical volumes for open system computers. Nine copies of a logical volume may be maintained within the same system.

Hitachi TrueCopy™

Hitachi TrueCopy basic software provides synchronous remote copy capability for open systems and S/390 computers. This allows remote copies over virtually unlimited distances. Operating systems that are supported include IBM MVS®, HP/UX®, IBM AIX®, Sun™ Solaris™, Digital® UNIX®, IBM (Sequent®) DYNIX/ptx®, SGI™ IRIX®, NCR®, UNIX® SVR4®, and Microsoft® Windows NT®/Windows® 2000.

Hitachi TrueCopy asynchronous software provides asynchronous remote copy capability for data from one Lightning 9900 V Series system to another Lightning 9900 V Series system or over channel extenders available from McDATA®, CNT® and INRANGE®. Records are timestamped to ensure that data is not lost or out of order. Since the transmission is asynchronous, data can be transmitted over unlimited distances. See also Chapter 3 for a discussion of the family of Hitachi TrueCopy software utilities.

Host Storage Domains

Host Storage Domains are collections of LUNs assigned to a specific host type in the Lightning 9900 V Series. Each Domain has its own logical (or virtual) Fibre Channel port and hosts, which are matched to their assigned Domain based on their WWNs. See also LUN, Fibre Channel, and WWNs.

Hot site

An alternate facility that has the equipment and resources to recover the business functions affected by the occurrence of a disaster or business interruption. Hot sites may vary in type of facilities offered (such as data processing, communication, or any other critical business functions needing duplication). Location and size of the hot site will be proportional to the equipment and resources needed. Similar terms include backup site, recovery site, recovery center, and alternate processing site. See also Cold site, Warm site, Disaster Recovery, Business Interruption, and Business Continuity Planning.

HXRC (Hitachi Extended Remote Copy)

This IBM® XRC™-compatible host-based software is offered by Hitachi Data Systems for asynchronous remote copy in an IBM S/390 environment. Using System Data Mover software, it guarantees data integrity for dependent write applications.

HYPERtape

HYPERtape is an enterprise backup/restore solution that leverages current customer investments. HYPERtape is a three-tier distributed system architecture with central administration and control that supports consolidated and distributed environments. HYPERtape can be used to back up data from any supported host to any system that supports the ftp protocol. Back up to disk for HSM integration or back up to local or network attached tape. Over 30 host platforms are covered and 70 library modules are supported. All popular RDBMS programs are supported, including Oracle®, SAP® R/3™, Informix®, Sybase®, DB2™, Adabas-D™ RDB, Microsoft® SQL Server®, Exchange, and Windows NT®/Windows® 2000 registry.

I/O

Input/output.

IP

The IP (Internet Protocol) is the underlying protocol for routing packets on the Internet and other TCP/IP-based networks. IP is an internetwork protocol that provides a communication standard that works across different types of linked networks (for example, Ethernet, FDDI, or ATM).

LAN

Local area networks or LANs are networks of computers that are geographically close together; this usually means on the same campus. Most LANs are confined to a single building or group of buildings. However, one LAN can be connected to other LANs over any distance via telephone lines, high-speed fiber optic backbones, and radio waves. A system of LANs connected in this way is called a wide-area network (WAN).

Lightning 9900™ V Series

The Hitachi Freedom Storage™ Lightning 9900™ V Series was announced in May 2002 (Lightning 9980V™ and Lightning 9970V™). These products represent a major advance in enterprise-class storage systems. This is because the second generation Hierarchical Star Network switched internal architecture provides for many times more simultaneous transfers to and from the host compared to shared-bus architectures. (The Lightning 9960™ and Lightning 9910™ systems of the Hitachi Freedom Storage Lightning 9900 Series were announced in June and November 2000 respectively.)

Logical Fibre Channel Port

The Lightning 9900 V Series Fibre Channel adapter cards each contain four physical Fibre Channel Ports. Each physical port can be subdivided into two or more logical Fibre Channel ports. Each logical Fibre Channel port supports a Host Storage Domain (HSD), which is a collection of LUNs matched to assigned hosts based on WWNs.

Logical Unit

The SCSI term for a logical disk drive.

Logical Unit Number

See LUN.

Logical Volume

The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. A host system sees a logical volume as a physical volume, although it does not correlate directly with a physical disk drive.

LUN

Logical Unit Number, or LUN, is a SCSI term for the field in an identifying message that is used to select a logical unit on a given target.

LPAR

LPAR or logical partition is an IBM® ESA/390® term for a set of functions that create the programming environment that is defined by the ESA/390 architecture. ESA/390 architecture uses this term when more than one LPAR is established on a an ESA/390 server. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is tied to one or more physical processors in a tightly coupled multiprocessor system. Also the LPAR does not depend on an operating system to create the virtual machine environment.

LUSE

The LUN Size Expansion feature of the Hitachi Freedom Storage™ allows standard size LUNs to be combined to create larger LUNs.

MAN

Metropolitan Area Networks are networks within a metropolitan area, which might, for example, be used for a city government.

MIB

Management Information Base is a set of standards for detailed system information that is reported to a control console for SNMP compliance. Its intent is to provide common parameters for heterogeneous computer systems.

MIPS

Millions of Instructions Per Second (or MIPS) is a rough measure of processor performance within the same class of processor.

Mirrored pair

Two disk units or logical units that contain the same data. The operating system software refers to them as one entity and "reads from either" and "writes to both" when RAID-1 is enabled.

Mirroring

A term to describe the process of writing data to two disk volumes, usually to ensure high availability in case one of the disks fails. Mirroring can be hardware or software based.

MPLF

The Hitachi Freedom Storage™ Lightning 9900™ V Series supports the Multiple-path Locking Facility (MPLF) for the IBM highest performance transaction processing operating system—TPF. In either native TPF mode or under VM, MPLF provides extremely high performance record-level locking so that multiple hosts can read and write to the same file without interfering with each other. See also TPF.

MTBF

Mean Time Between Failure. A commonly used measure of system reliability, usually expressed in hours. Modern disk drives typically have an MTTR of 1 million hours or more.

MTTR

Mean Time To Repair. Includes the time taken to diagnose the failure, replace, or repair faulty component(s) and restart the system so it is available to users. See MTBF.

NanoCopy™

NanoCopy is a feature of the Lightning 9900 Series and Lightning 9900 V Series product lines that enables time-consistent snapshots to be taken without stopping applications to flush in-flight data to disk. Since there is no system impact in taking a NanoCopy snapshot, snapshots can be made more frequently for faster recovery in the event of a failure. See also ShadowImage.

NAS

Network Attached Storage or NAS servers are a special class of server that allows files to be stored over networks using the UNIX® or Microsoft® Windows® remote file system standards.

N+1

An N+1 power supply design provides for one redundant power supply in a power system design that provides full system power in the event of a power supply failure.

NDMP

Network Data Management Protocol (NDMP) is a standard protocol for network-based backup of network-attached storage. NDMP hides the unique hardware interfaces from third-party backup software that allows this software to execute on any NDMP compliant system on the network.

Node

See Fibre Channel.

Off-site storage facility

A secure location, remote from the primary location, at which backup hardware, software, data files, documents, equipment, or supplies are stored.

Online systems

An interactive computer system supporting users over a network of computer terminals.

Open system

A system whose characteristics comply with de facto standards made available throughout the industry, and therefore can be connected to other systems that comply with the same standards.

Operating system

The operating system is the most important software program that runs on a computer. The operating system (OS) performs basic tasks such as recognizing input from a keyboard, sending output to the display screen, keeping track of files and directories on the disk, and controlling peripheral devices such as disk drive and printers or a mouse. The OS acts as a traffic cop and schedules the various programs that the computer executes. The OS is also responsible for security, ensuring that unauthorized users do not access the system. Operating systems can be classified as follows:

- 1) Multi-user—allows two or more users to run programs at the same time.
- 2) Multi-processing—supports running a program on more than one CPU.
- 3) Multi-tasking—allows more than one program to run concurrently.
- 4) Multi-threading—allows different parts of a single program to run concurrently.
- 5) Real Time—Usually a stripped down OS that responds to input instantly.

PiT

A point-in-time (PiT) copy is a copy of data that is taken at a specific point in time. PiT copies are used in many ways, including backups and checkpoints.

Port/Port ID

See Fibre Channel.

PPRC

Peer-to-Peer Remote Copy or PPRC is an IBM feature that maintains consistent copy of a logical volume on the same controller or on another controller. Access to the remote controller occurs over an ESCON path. All writes that any attached host performs on the primary logical volume, are also performed on the secondary logical volume. The user can request that the controller perform the modifications synchronously, that is, before the controller completes the modification of the primary copy.

QoS

Quality of service refers to service level agreement (SLAs) between users and the IT organization. See also SLA.

RAID

Redundant Array of Independent Disks. RAID is used to increase the reliability of disk arrays by providing redundancy either through complete duplication of the data (RAID-1, i.e., mirroring) or through construction of parity data for each data stripe in the array (RAID-3, -4, -5). RAID-5, which distributes parity information across all disks in an array, is among the most popular means of providing parity RAID since it avoids the bottlenecks of a single parity disk. The Hitachi Freedom Storage™ Lightning 9900™ V Series system algorithms enable performance from RAID-5 that is competitive with some vendors' RAID-1. Some vendors do not offer RAID-5.

RAID controllers

RAID controllers provide a highly optimized scheme for securely managing RAID configurations on storage systems. Hitachi RAID controllers allow RAID arrays to be expanded online, and support conversion of an array from one RAID level to another.

Recovery time

The period from the onset of a disaster to the recovery of the critical functions.

Remote copy

Remote copy refers generically to software or hardware utilities that provide the capability to copy data from one online volume to remote volumes without disruption. Synchronous techniques are used for short distances (typically less than 25 miles) and asynchronous techniques over LAN/WAN/MAN are used at any distance. For a complete discussion of Hitachi TrueCopy™ (remote copy) solutions, see Chapter 3.

Remote copy links

This term refers to the links used between storage systems for the movement of data. Today these links are either direct connect IBM® ESCON®, Fibre Channel, or network links (T3, ATM, etc). For direct connect ESCON there is a limit of 43km (25 miles). For direct connect Fibre Channel the limit is 10km, however newer technologies, such as the Nortel® OPTERA™ product are allowing direct fibre connect over longer distances.

Risk management

The discipline that ensures that an organization does not assume an unacceptable level of risk.

SAN

Storage area networks (SANs) connect storage systems to servers through Fibre Channel or Ethernet switches. The Hitachi implementation of the SAN is known as Hitachi Freedom Data Networks™. Major benefits of SANs include outboard backup, sharing of resources, pooling, and reduced cost of storage management. SANs are also defined as high-speed subnetworks of shared storage devices. SAN architecture works in a way that makes all storage devices available to all servers on a LAN or WAN. Because stored data does not reside directly on any of a network's servers, server power is utilized for business applications, and network capacity is released to the end user. See also Hitachi Freedom Data Networks and Fibre Channel.

SCSI

Small Computer System Interface. An intelligent bus-level interface that defines a standard I/O bus and a set of high-level I/O commands. There are currently many flavors of SCSI defined by different bus widths and clock speeds. The seven major variations of SCSI are SCSI 1, SCSI 2 (Fast/Narrow), SCSI 2 (Fast/Wide), Ultra SCSI (Fast/Narrow), Ultra SCSI (Fast/Wide) – also called SCSI 3, Ultra 2 SCSI (Narrow), Ultra 2 SCSI Wide. See also Fibre Channel.

Serverless Backup

Using the eXtended SCSI command and products like VERITAS® NetBackup™, the Hitachi Freedom Storage™ Lightning 9900™ V Series offers serverless and LANfree backup.

ShadowImage LVD

ShadowImage LVD (Logical Volume Divider) is a S/390 product that controls the execution of the ShadowImage function of the Lightning 9900 or Lightning 9900V disk array subsystems.

SIM

Service Information Messages are messages generated by Lightning 9900 Series, Lightning 9900 V Series, and Thunder 9200™ systems to identify normal operations.

SLA

Service Level Agreements or SLAs are agreements regarding level of service between user departments and the IT department. SLAs refer to all aspects of IT service, including availability, performance, and repair.

Snapshot

A term that refers to a copy of a file system at a certain point in time. Snapshots are used for backup and recovery.

SNMP

Simple Network Management Protocol. SNMP is a protocol used for communication between simple, server-resident SNMP agents that respond to network administration requests from simple-to-sophisticated SNMP manager tools running on remote workstations.

Stripe

In RAID terminology, a stripe is when data is read or written in parallel to or from multiple disks instead of reading or writing all data to one disk. Striping provides much higher performance through its parallel design.

SWAN

Storage wide area networks (SWANs) are interconnected SANs over long distances. They are made possible by Fibre Channel and IBM® ESCON® extenders.

Switch

A switch is a network device that examines and forwards packets between LAN segments.

Synchronous

Synchronous communications occur when the transmission of data between two devices is synchronized with a clocking scheme or other technique. The sender and receiver need to synchronize with one another before data is sent. In synchronous communication, the bit stream and the clock pulse are synchronized by a special bit transition pattern in the digital signal, creating an exactly timed stream of bits from the sending device to the receiving device. An example of such a mechanism is bipolar encoding. Synchronous communication is either character or bit oriented. Character-oriented synchronous transmissions are used to send blocks of characters such as found in ASCII (American Standard Code for Information Interchange) files. Bit-oriented synchronous communication is used primarily for the transmission of binary data. See also Asynchronous, Binary, and Hitachi TrueCopy™.

TCO

Total cost of ownership or TCO is a computer industry financial method of identifying the cost of operating computer equipment. TCO includes depreciated capital costs, manpower expense, power costs, communication costs, overhead, etc.

TCP

Transmission Control Protocol or TCP is a transport layer component of the Internet's TCP/IP protocol suite. It sits above IP in the protocol stack and provides reliable data delivery services over connection-oriented links. TCP uses IP to deliver information across a network and makes up for the deficiency of IP providing a guarantee of reliable delivery services that IP does not. TCP messages and data are encapsulated into IP datagrams, which IP then delivers across the network.

TPF

Transaction Processing Facility or TPF is highest performance transaction processing software environment from IBM. TPF is used by many of the world's largest Customer Reservation Systems such as SABRE® and by the world's largest financial institutions. TPF uses a high performance record-locking scheme called MPLF. See also MPLF.

VERITAS®

A Mountain View, California, software company that develops and supports volume and file management software products for a variety of UNIX® and Microsoft® Windows® platforms.

Virtual Logical Volume Image Manager

Virtual Logical Volume Image (VLVI) Manager is a software utility in the Hitachi Resource Manager™ suite that allows for configuration of RAID, as well as create, delete, verify, rebuild, tune, and abort operations. See also RAID.

Volume

An IBM® ESA/390® term for the information recorded on a single disk unit or recording medium. Indirectly, a volume can refer to the unit of recording medium itself. On a non-removable medium storage device such as a disk drive, the terms may also refer, indirectly, to the storage device that is associated with the volume. When a user stores multiple volumes on a single storage medium transparent to the program, the volumes are referred to as logical volumes.

WAN

Wide area networks or WANs are networks of computers that are geographically dispersed and connected by radio waves, telephone lines, satellites, or high-speed fiber optic backbones.

Warm site

An alternate-processing site that is only partially equipped (as compared to hot site, which is fully equipped). See also Hot site and Cold site.

WWNs

World Wide Names (WWNs) refer to an eight-byte identifier assigned to each product that can be used as a port on a Fibre Channel network. The WWN is stored in nonvolatile memory and is frequently stamped on the surface of the product or used as a serial number. It applies to all HBAs, switches, or storage controller cards that interface to a Fibre Channel network.

XRC™

Extended Remote Copy is an IBM implementation of a software asynchronous remote copy technique that preserves data integrity. See also Chapter 3 for a complete discussion of remote copy software and Hitachi TrueCopy™ product family.

Hitachi Data Systems Worldwide Offices

www.hds.com

www.storage@hds.com

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia Headquarters

Suite 3301-6, Shell Tower
Times Square, 1 Matheson Street
Causeway Bay
Hong Kong
2525-2385
infoasia@hds.com

Australia/New Zealand Headquarters

Hitachi Data Systems
Level 3
82 Waterloo Rd.
North Ryde, NSW 2113
Australia
+ 61-2-9325-3300
info.australia@hds.com

Canada Headquarters

2550 Victoria Park Avenue
Suite 601
Toronto, Ontario M2J 5A9
Canada
(416) 494-4114
www.hds.com
Canada.Sales@hds.com

Europe Headquarters

Sefron Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
+44 (0) 1753-618000
info.eu@hds.com
www.eu.hds.com

Latin America Headquarters

750 Central Expressway, MS 3268
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
infolatin@hds.com

U.S. Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
ussalesinfo@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and a service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Hi-Track is a registered service mark and Freedom Storage, Lightning 9900, Thunder 9200, HiCommand, Resource Manager, ShadowImage, Dynamic Link Manager, RapidXchange, Freedom Data Networks, SANtinel, Graph-Track, FlashAccess, NanoCopy, Thunder 9500, Lightning 9970V, Lightning 9980V, Thunder 9570V, Thunder 9531V, Thunder 9532V, Thunder 9533V, CruiseControl, Lightning 9960, Lightning 9910, TrueCopy, TrueNorth, and Hi-Star are trademarks of Hitachi Data Systems Corporation.

All other trademarks, trade names, and service marks used herein are the property of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to <http://www.hds.com/products/software/licenses.html> or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

©2003, Hitachi Data Systems Corporation
DISK-436-02 January 2003